

EXHIBIT 5

DocCode – SCORE

SCORE Placeholder Sheet for IFW Content

Application Number: 15061021

Document Date: 03/04/2016

The presence of this form in the IFW record indicates that the following document type was received in electronic format on the date identified above. This content is stored in the SCORE database.

Since this was an electronic submission, there is no physical artifact folder, no artifact folder is recorded in PALM, and no paper documents or physical media exist. The TIFF images in the IFW record were created from the original documents that are stored in SCORE.

- Drawing

At the time of document entry (noted above):

- USPTO employees may access SCORE content via eDAN using the Supplemental Content tab, or via the SCORE web page.
- External customers may access SCORE content via PAIR using the Supplemental Content tab.

Doc Code: PA..

6165

PTO/AIA/82A (07-12)

Document Description: Power of Attorney

Approved for use through 11/30/2014. OMB 0651-0035

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE
REGISTERED PRACTITIONERS**

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B or equivalent) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5. If the Power of Attorney by Applicant form is not accompanied by this transmittal form or an equivalent, the Power of Attorney will not be recognized in the application.

Application Number	
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Title	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
Art Unit	
Examiner Name	
Attorney Docket Number	YOR920120450US2

SIGNATURE of Applicant or Patent Practitioner

Signature	/Paul J. Otterstedt/	Date	2016-03-04
Name	Paul J. Otterstedt	Telephone	(203) 433-3157
Registration Number	37411		

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications.

*Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32 and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☒ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above:

98841

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☒ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

Firm or Individual Name			
Address			
City	State	Zip	
Country			
Telephone	Email		

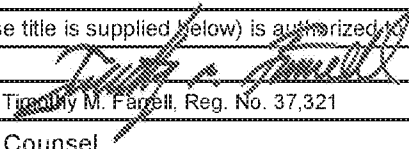
I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

International Business Machines Corporation

- ☐ Inventor or Joint Inventor (title not required below)
- ☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)
- ☒ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)
- ☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE OF Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature		Date (Optional)	3/10/15
Name	Timothy M. Farrell, Reg. No. 37,321		
Title	Counsel		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☒ Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		
<p>The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.</p> <p>This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.</p>			

Secrecy Order 37 CFR 5.2:

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor	1				Remove
Legal Name					
Prefix	Given Name	Middle Name	Family Name	Suffix	
	Matthew	A.	Markley		
Residence Information (Select One) <input checked="" type="radio"/> US Residency <input type="radio"/> Non US Residency <input type="radio"/> Active US Military Service					
City	Hays	State/Province	KS	Country of Residence	US

Mailing Address of Inventor:

Address 1	1582 280th Avenue				
Address 2					
City	Hays	State/Province	KS		
Postal Code	67601	Country	US		

Inventor	2				Remove
Legal Name					
Prefix	Given Name	Middle Name	Family Name	Suffix	
	Amitkumar	M.	Paradkar		
Residence Information (Select One) <input checked="" type="radio"/> US Residency <input type="radio"/> Non US Residency <input type="radio"/> Active US Military Service					
City	Mohegan Lake	State/Province	NY	Country of Residence	US

Mailing Address of Inventor:

Address 1	BM T.J. Watson Research Center				
Address 2	1101 Kitchawan Rd				
City	Yorktown Heights	State/Province	NY		
Postal Code	10598	Country	US		

Inventor	3				Remove
Legal Name					
Prefix	Given Name	Middle Name	Family Name	Suffix	
	Venkata	Vinay	Parisa		
Residence Information (Select One) <input checked="" type="radio"/> US Residency <input type="radio"/> Non US Residency <input type="radio"/> Active US Military Service					

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

City	Charlotte	State/Province	NC	Country of Residence	US
------	-----------	----------------	----	----------------------	----

Mailing Address of Inventor:

Address 1	10330 David Taylor Drive,2D				
Address 2					
City	Charlotte	State/Province	NC		
Postal Code	28262	Country i	US		
Inventor	4				<input type="button" value="Remove"/>
Legal Name					

Prefix	Given Name	Middle Name	Family Name	Suffix
	Birgit	M.	Pfitzmann	
Residence Information (Select One) <input type="radio"/> US Residency <input checked="" type="radio"/> Non US Residency <input type="radio"/> Active US Military Service				

City	Samstagern	Country of Residence i	CH
------	------------	------------------------	----

Mailing Address of Inventor:

Address 1	Bärenweidstrasse 23				
Address 2					
City	Samstagern	State/Province			
Postal Code	CH-8833	Country i	CH		
All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the Add button.					

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below. For further information see 37 CFR 1.33(a).			
<input type="checkbox"/> An Address is being provided for the correspondence Information of this application.			
Customer Number	98841		
Email Address	DOCKET@OEKLAW.COM	<input type="button" value="Add Email"/>	<input type="button" value="Remove Email"/>

Application Information:

Title of the Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		
Attorney Docket Number	YOR920120450US2	Small Entity Status Claimed	<input type="checkbox"/>
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	47	Suggested Figure for Publication (if any)	16

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

Filing By Reference:

Only complete this section when filing an application by reference under 35 U.S.C. 111(c) and 37 CFR 1.57(a). Do not complete this section if application papers including a specification and any drawings are being filed. Any domestic benefit or foreign priority information must be provided in the appropriate section(s) below (i.e., "Domestic Benefit/National Stage Information" and "Foreign Priority Information").

For the purposes of a filing date under 37 CFR 1.53(b), the description and any drawings of the present application are replaced by this reference to the previously filed application, subject to conditions and requirements of 37 CFR 1.57(a).

Application number of the previously filed application	Filing date (YYYY-MM-DD)	Intellectual Property Authority or Country

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)

☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not be** the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32). Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	<input checked="" type="radio"/> Customer Number	US Patent Practitioner	<input type="radio"/> Limited Recognition (37 CFR 11.9)
Customer Number	98841		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, 365(c), or 386(c) or indicate National Stage entry from a PCT application. Providing benefit claim information in the Application Data Sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

When referring to the current application, please leave the "Application Number" field blank.

Prior Application Status	Pending	Remove	
Application Number	Continuity Type	Prior Application Number	Filing or 371(c) Date (YYYY-MM-DD)
	Continuation of	13689500	2012-11-29
Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the Add button.			Add

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55. When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)ⁱ the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(i)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Application Number	Country ⁱ	Filing Date (YYYY-MM-DD)	Access Code ⁱ (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.

NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

Authorization or Opt-Out of Authorization to Permit Access:

When this Application Data Sheet is properly signed and filed with the application, applicant has provided written authority to permit a participating foreign intellectual property (IP) office access to the instant application-as-filed (see paragraph A in subsection 1 below) and the European Patent Office (EPO) access to any search results from the instant application (see paragraph B in subsection 1 below).

Should applicant choose not to provide an authorization identified in subsection 1 below, applicant **must opt-out** of the authorization by checking the corresponding box A or B or both in subsection 2 below.

NOTE: This section of the Application Data Sheet is **ONLY** reviewed and processed with the **INITIAL** filing of an application. After the initial filing of an application, an Application Data Sheet cannot be used to provide or rescind authorization for access by a foreign IP office(s). Instead, Form PTO/SB/39 or PTO/SB/69 must be used as appropriate.

1. Authorization to Permit Access by a Foreign Intellectual Property Office(s)

A. Priority Document Exchange (PDX) - Unless box A in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the State Intellectual Property Office of the People's Republic of China (SIPO), the World Intellectual Property Organization (WIPO), and any other foreign intellectual property office participating with the USPTO in a bilateral or multilateral priority document exchange agreement in which a foreign application claiming priority to the instant patent application is filed, access to: (1) the instant patent application-as-filed and its related bibliographic data, (2) any foreign or domestic application to which priority or benefit is claimed by the instant application and its related bibliographic data, and (3) the date of filing of this Authorization. See 37 CFR 1.14(h)(1).

B. Search Results from U.S. Application to EPO - Unless box B in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the EPO access to the bibliographic data and search results from the instant patent application when a European patent application claiming priority to the instant patent application is filed. See 37 CFR 1.14(h)(2).

The applicant is reminded that the EPO's Rule 141(1) EPC (European Patent Convention) requires applicants to submit a copy of search results from the instant application without delay in a European patent application that claims priority to the instant application.

2. Opt-Out of Authorizations to Permit Access by a Foreign Intellectual Property Office(s)

☐ A. Applicant **DOES NOT** authorize the USPTO to permit a participating foreign IP office access to the instant application-as-filed. If this box is checked, the USPTO will not be providing a participating foreign IP office with any documents and information identified in subsection 1A above.

☐ B. Applicant **DOES NOT** authorize the USPTO to transmit to the EPO any search results from the instant patent application. If this box is checked, the USPTO will not be providing the EPO with search results from the instant application.

NOTE: Once the application has published or is otherwise publicly available, the USPTO may provide access to the application in accordance with 37 CFR 1.14.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant	1	Remove
<p>If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.</p> <p style="text-align: right;">Clear</p>		
<input checked="" type="radio"/> Assignee	Legal Representative under 35 U.S.C. 117	Joint Inventor
Person to whom the inventor is obligated to assign.		Person who shows sufficient proprietary interest
If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:		
<div style="border: 1px solid black; height: 20px; width: 100%;"></div>		
Name of the Deceased or Legally Incapacitated Inventor: <div style="border: 1px solid black; height: 20px; width: 100%;"></div>		
If the Applicant is an Organization check here. <input checked="" type="checkbox"/>		
Organization Name	International Business Machines Corporation	
Mailing Address Information For Applicant:		
Address 1	New Orchard Road	
Address 2		
City	Armonk	State/Province NY
Country	US	Postal Code 10504
Phone Number		Fax Number
Email Address	IPLAWYOR@US.IBM.COM	
Additional Applicant Data may be generated within this form by selecting the Add button. Add		

Assignee Information including Non-Applicant Assignee Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

Assignee	1			
Complete this section if assignee information, including non-applicant assignee information, is desired to be included on the patent application publication. An assignee-applicant identified in the "Applicant Information" section will appear on the patent application publication as an applicant. For an assignee-applicant, complete this section only if identification as an assignee is also desired on the patent application publication.				
<input type="button" value="Remove"/>				
If the Assignee or Non-Applicant Assignee is an Organization check here. <input type="checkbox"/>				
Prefix	Given Name	Middle Name	Family Name	Suffix
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Mailing Address Information For Assignee including Non-Applicant Assignee:				
Address 1		<input type="text"/>		
Address 2		<input type="text"/>		
City	<input type="text"/>	State/Province	<input type="text"/>	
Country ⁱ	<input type="text"/>	Postal Code	<input type="text"/>	
Phone Number	<input type="text"/>	Fax Number	<input type="text"/>	
Email Address	<input type="text"/>			
Additional Assignee or Non-Applicant Assignee Data may be generated within this form by selecting the Add button. <input type="button" value="Add"/>				

Signature:				<input type="button" value="Remove"/>	
NOTE: This Application Data Sheet must be signed in accordance with 37 CFR 1.33(b). However, if this Application Data Sheet is submitted with the INITIAL filing of the application and either box A or B is not checked in subsection 2 of the "Authorization or Opt-Out of Authorization to Permit Access" section, then this form must also be signed in accordance with 37 CFR 1.14(c). This Application Data Sheet must be signed by a patent practitioner if one or more of the applicants is a juristic entity (e.g., corporation or association). If the applicant is two or more joint inventors, this form must be signed by a patent practitioner, all joint inventors who are the applicant, or one or more joint inventor-applicants who have been given power of attorney (e.g., see USPTO Form PTO/AIA/81) on behalf of all joint inventor-applicants. See 37 CFR 1.4(d) for the manner of making signatures and certifications.					
Signature	/PAUL J. OTTERSTEDT/			Date (YYYY-MM-DD)	2016-03-04
First Name	PAUL	Last Name	OTTERSTEDT	Registration Number	37411
Additional Signature may be generated within this form by selecting the Add button. <input type="button" value="Add"/>					

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	YOR920120450US2
		Application Number	
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

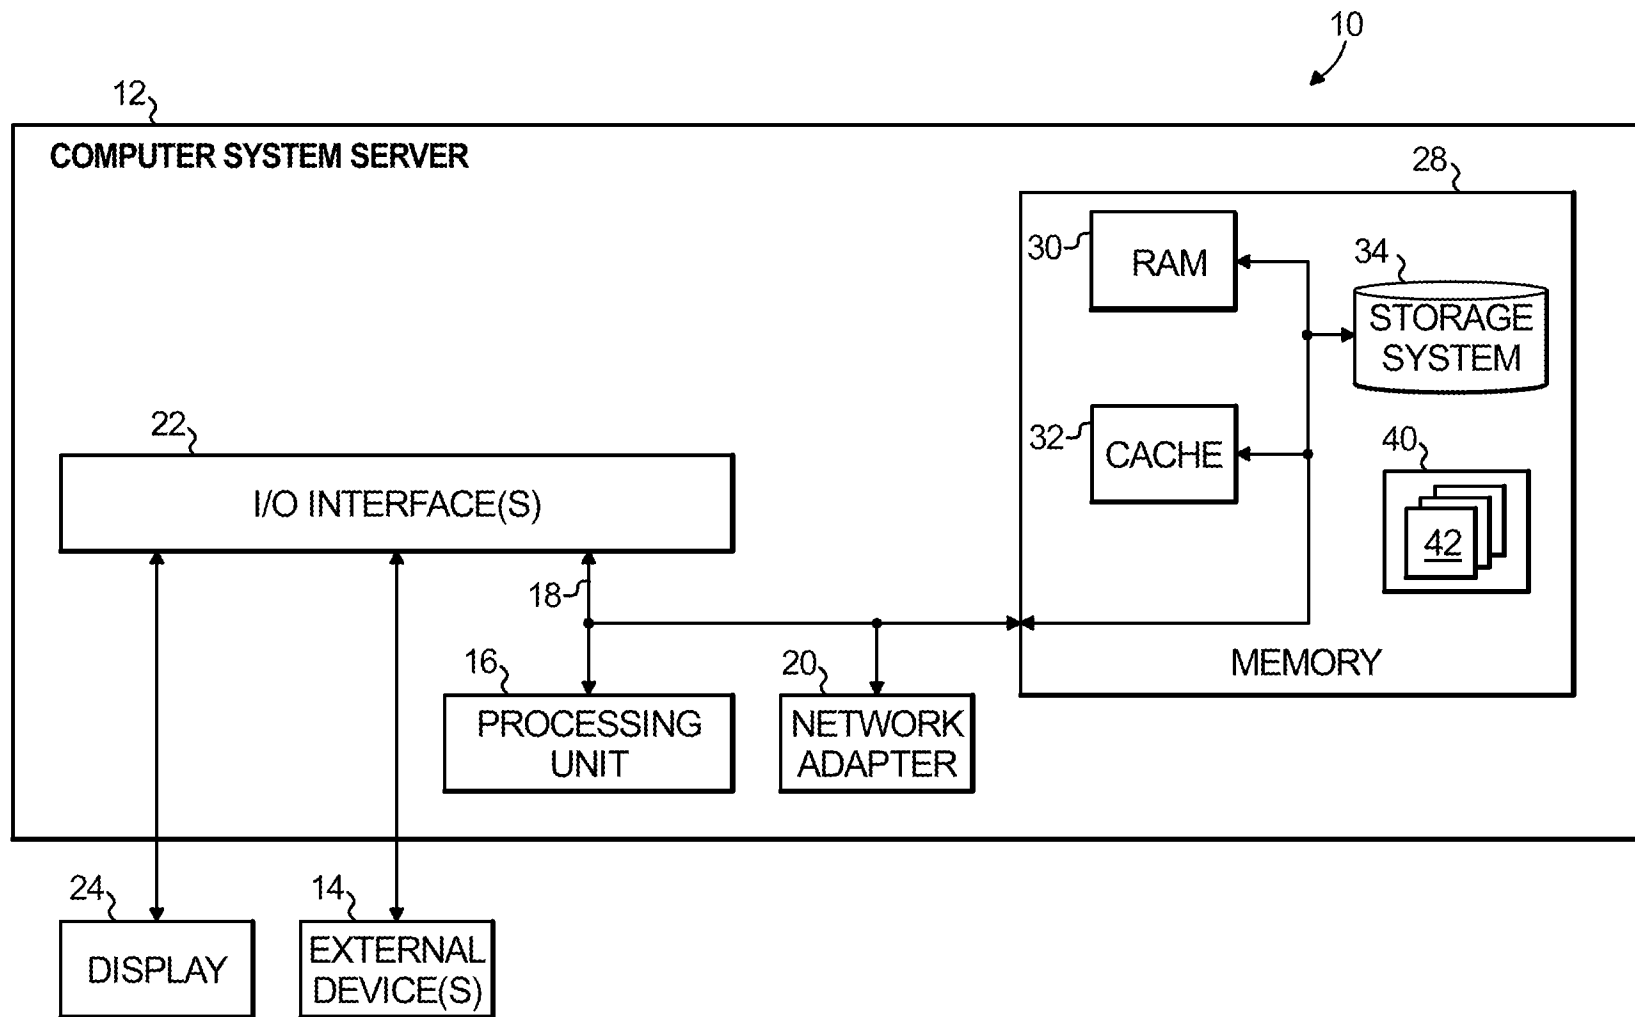


Fig. 1

2/47




Fig. 2

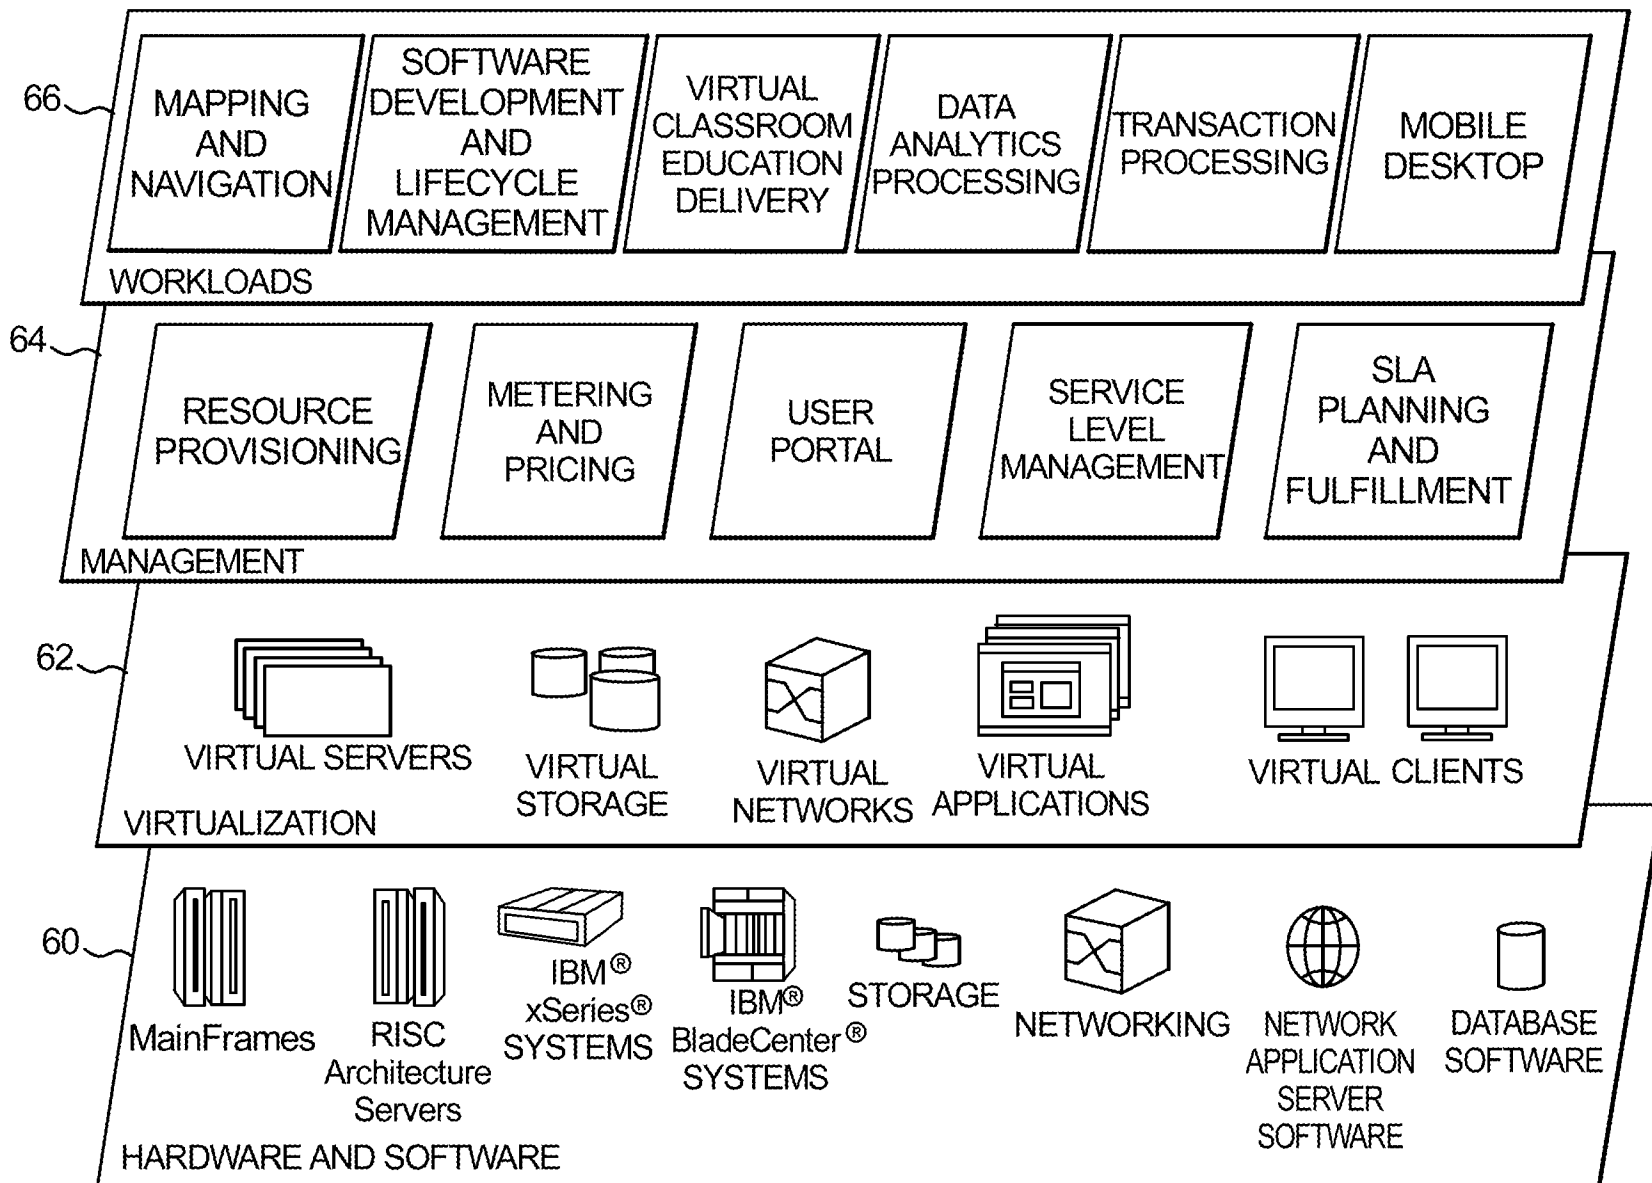


Fig. 3

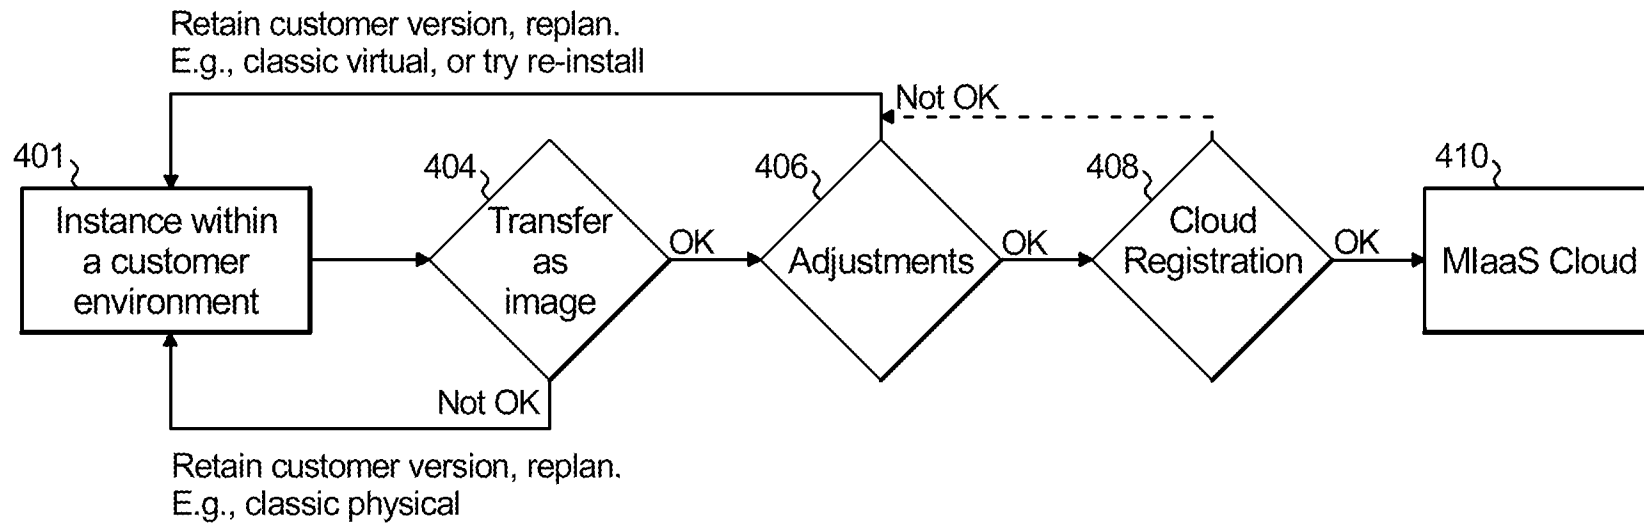


Fig. 4

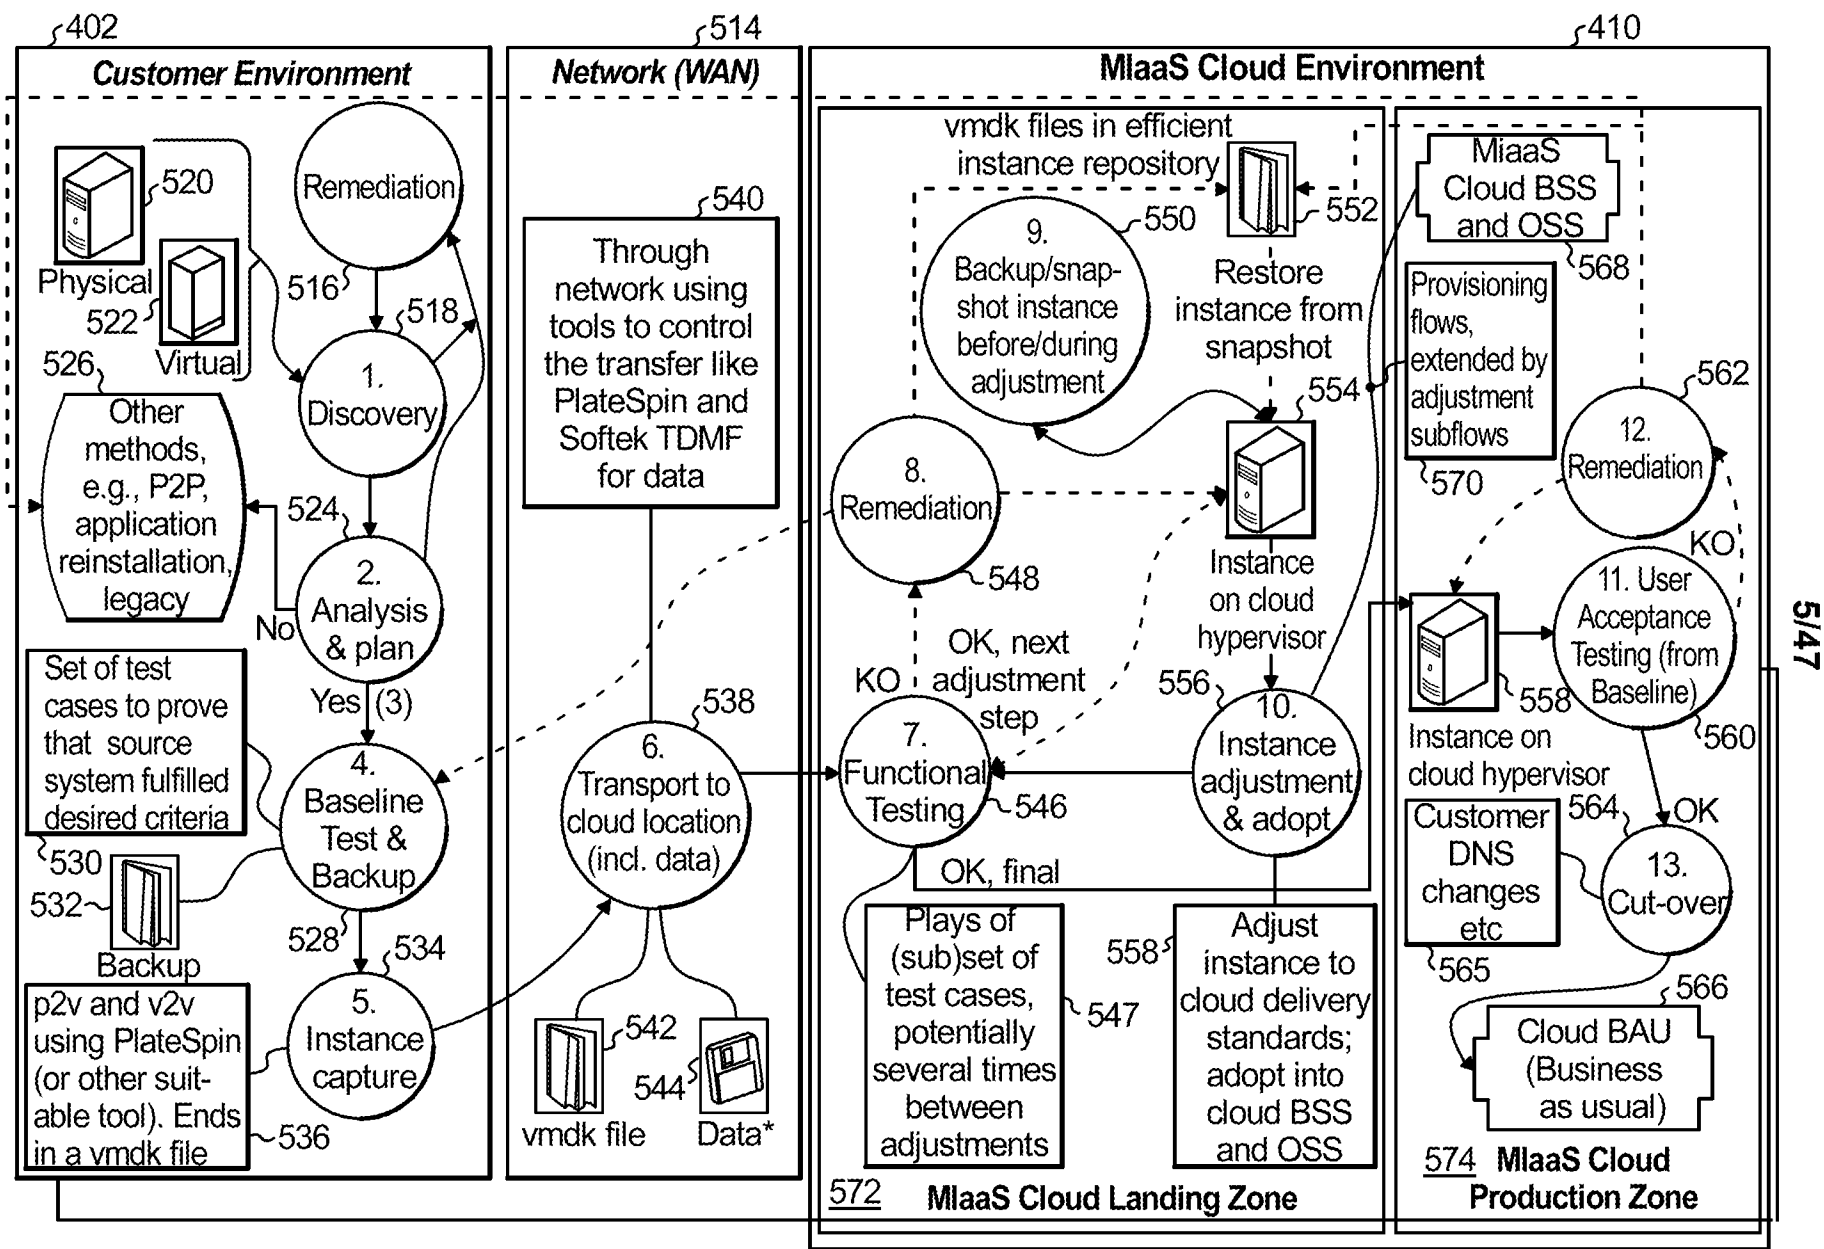


Fig. 5

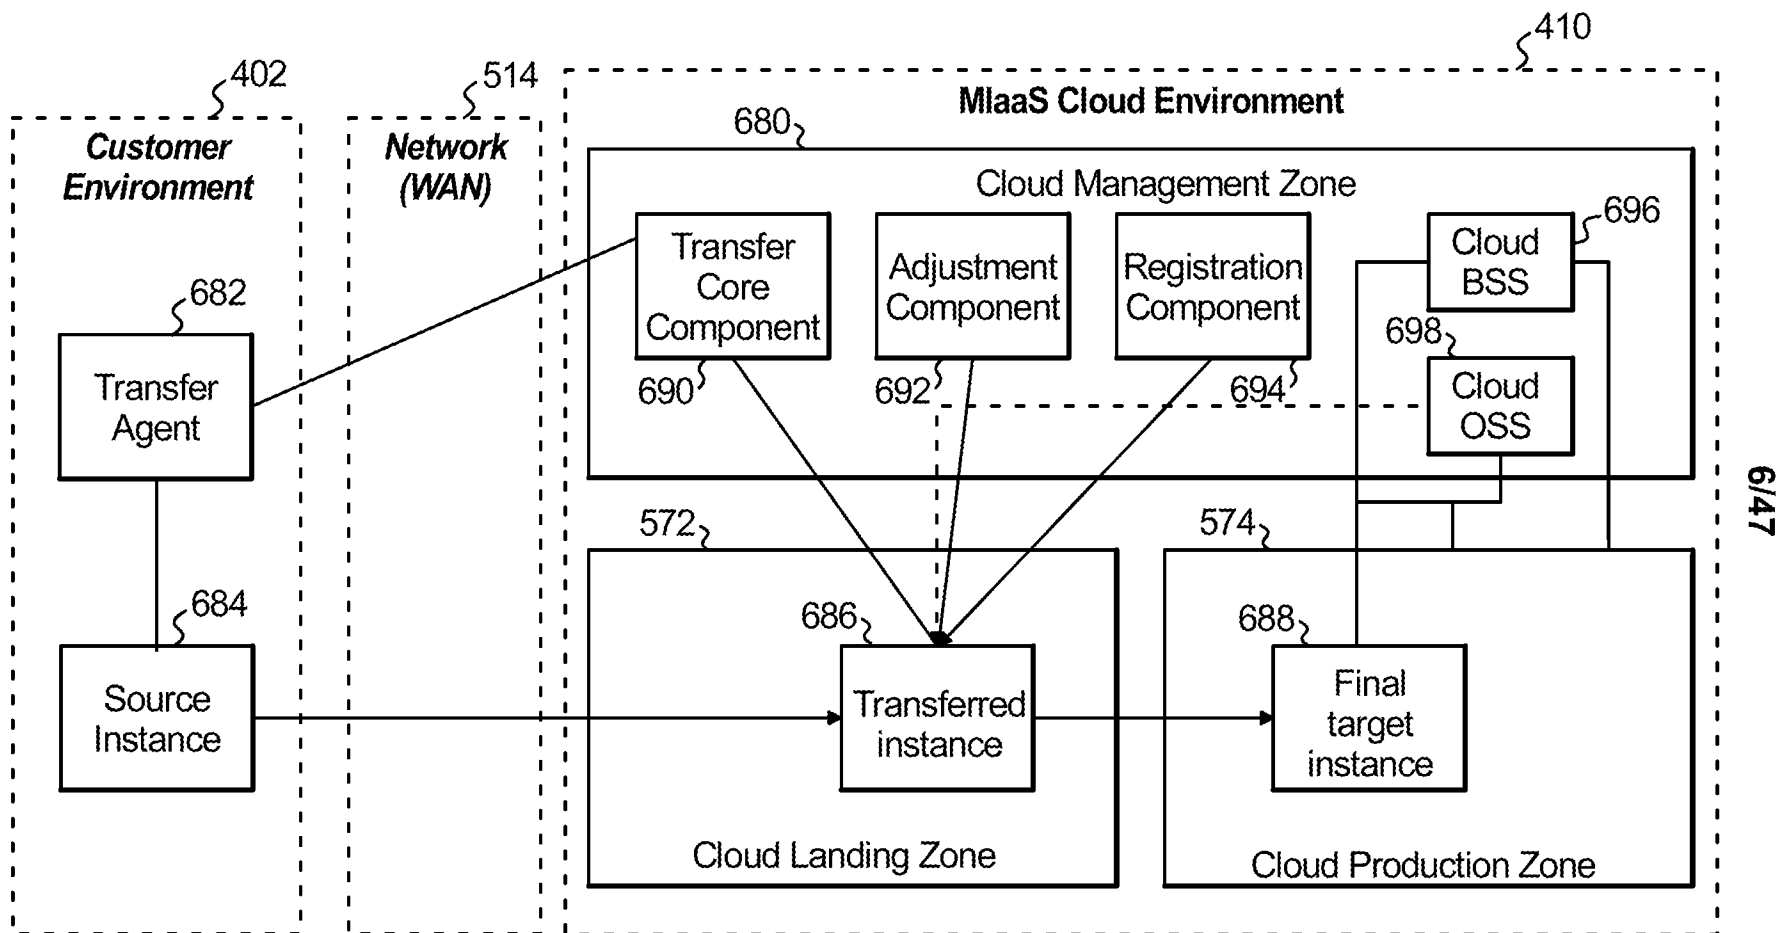



Fig. 6



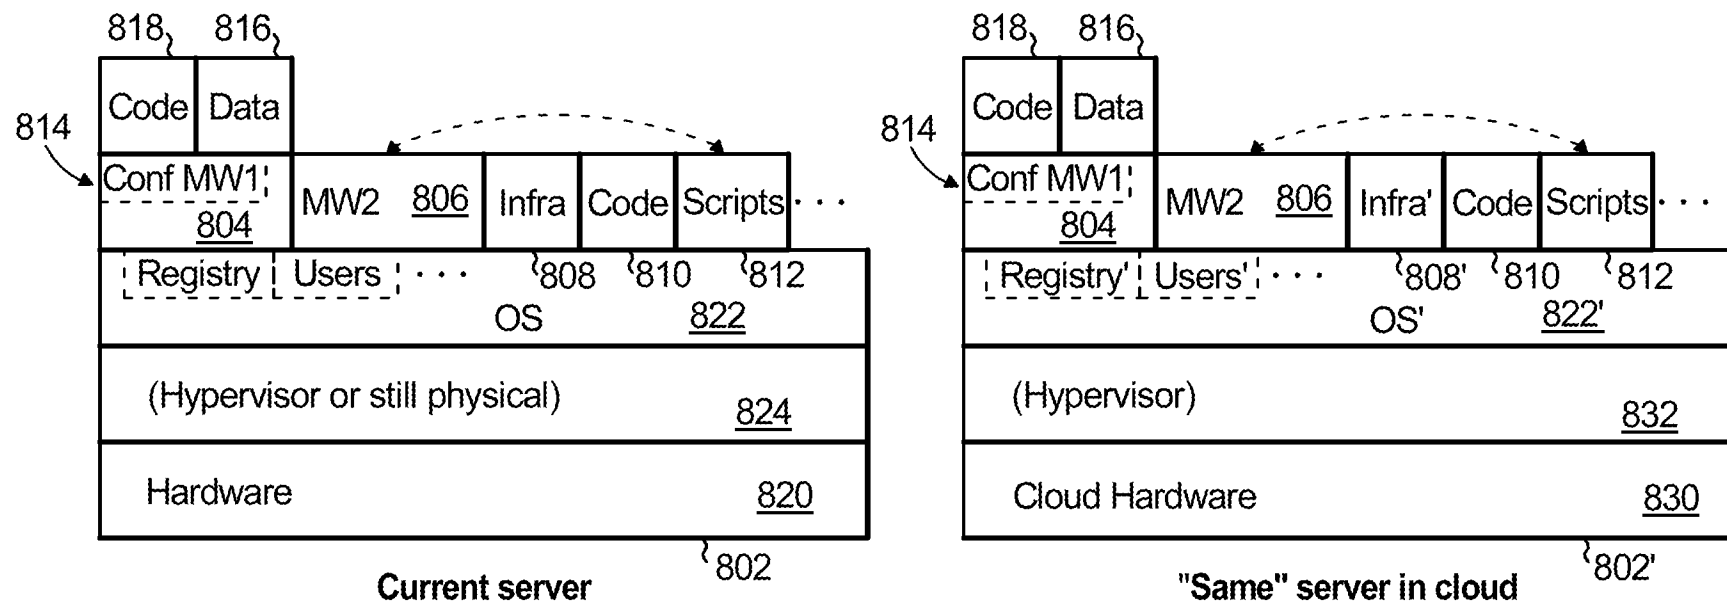


Fig. 8

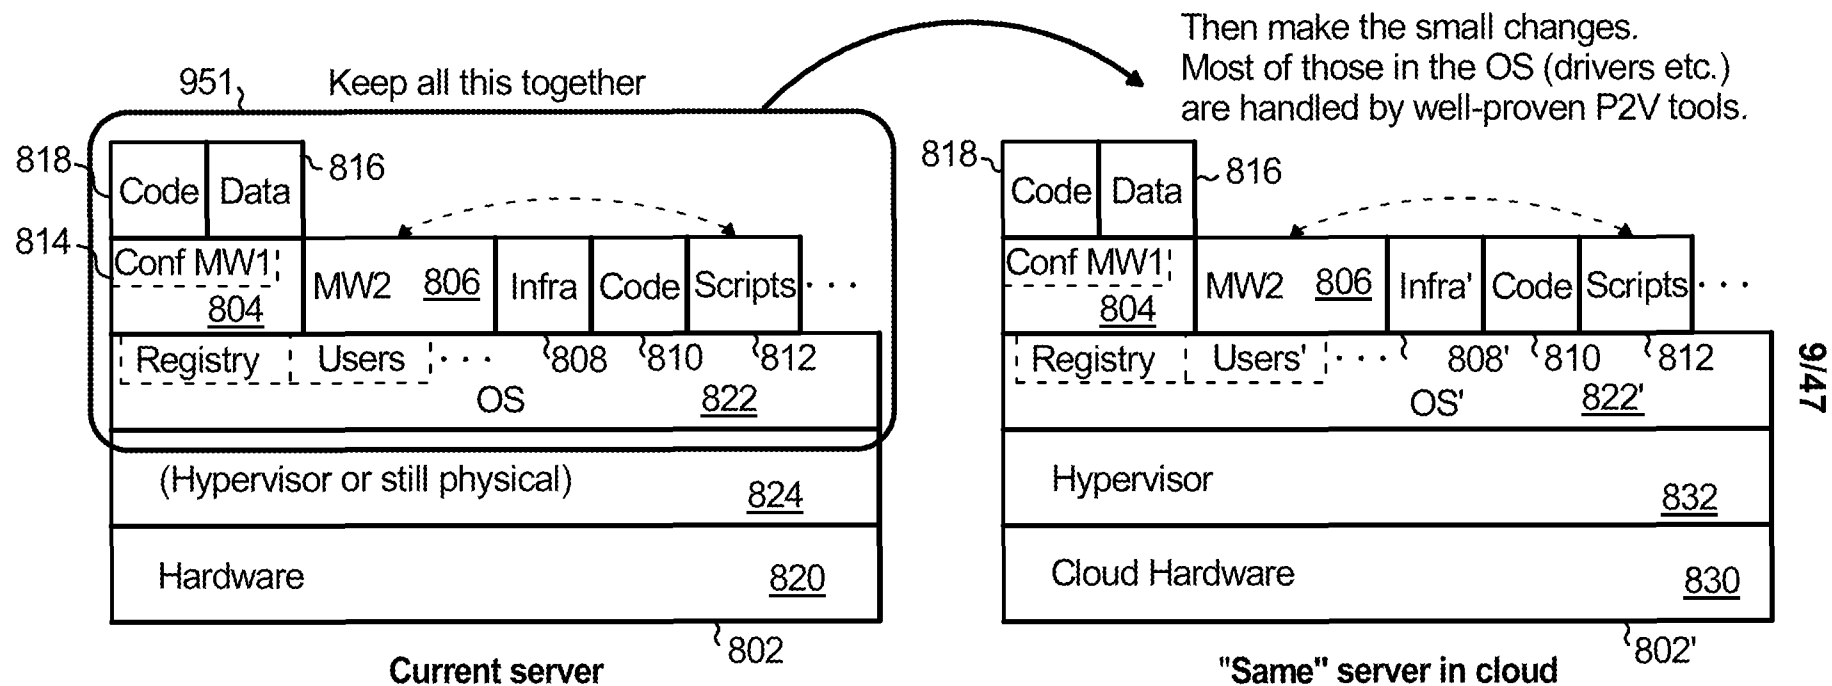


Fig. 9

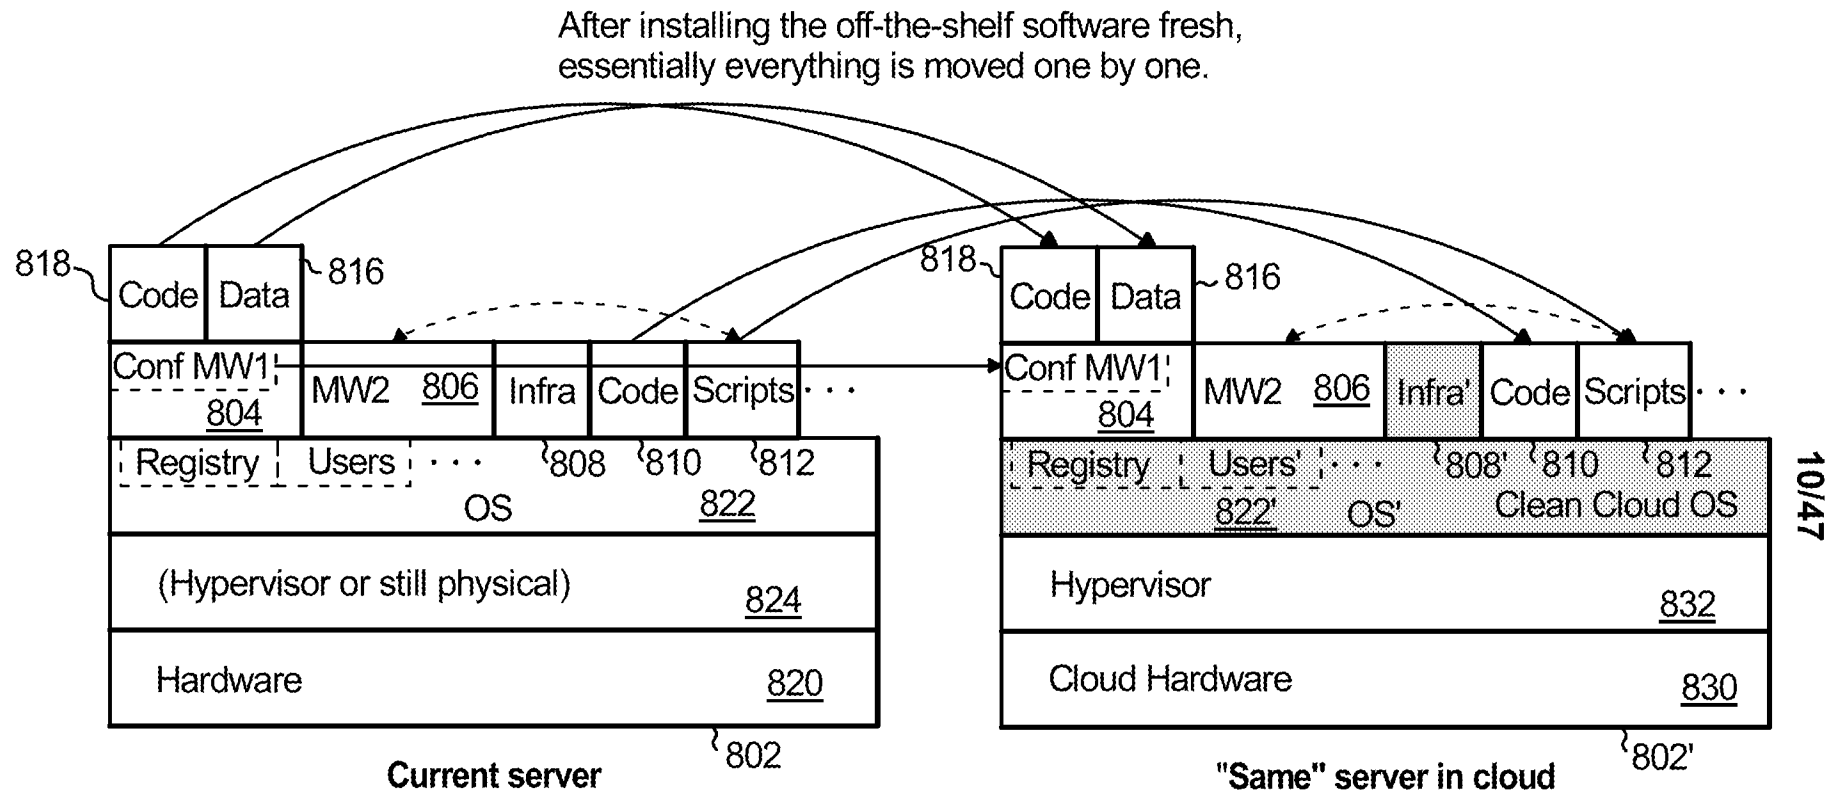
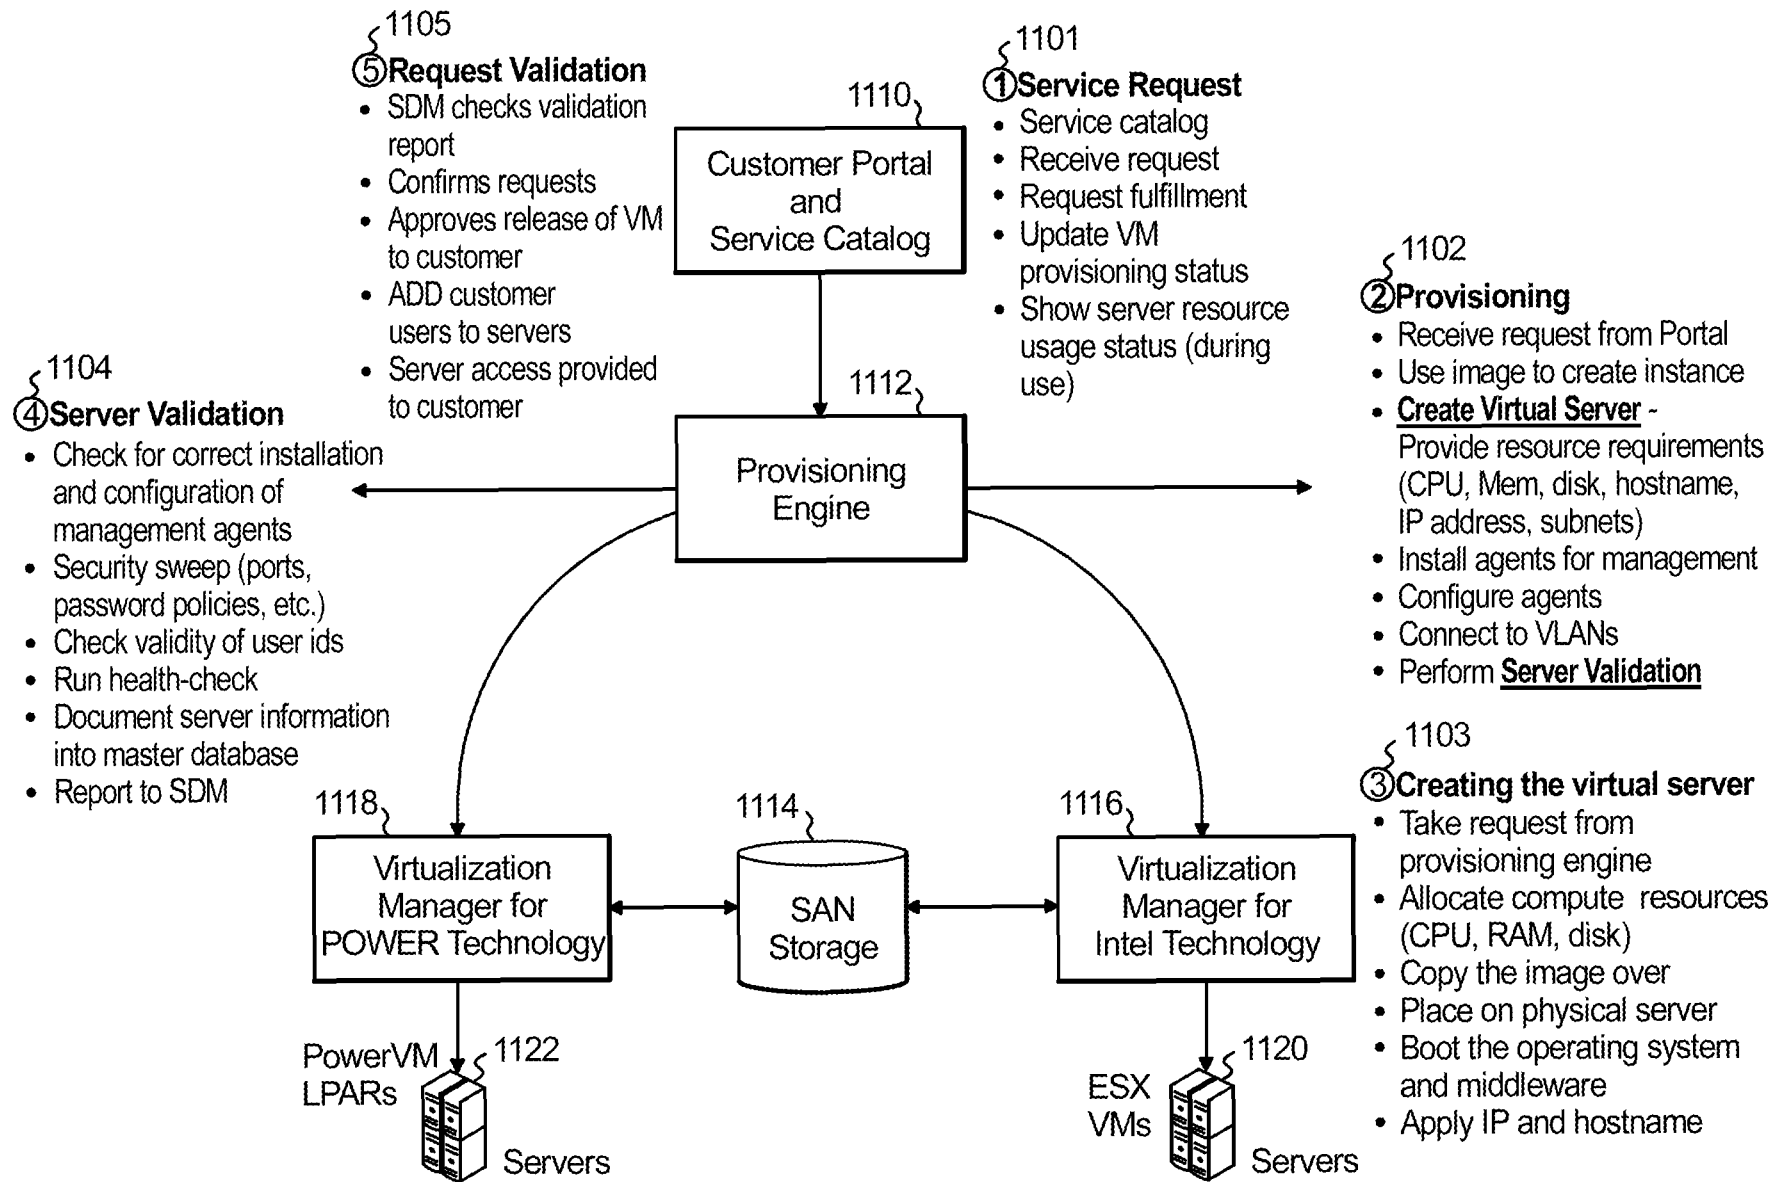


Fig. 10



11/47

YOR920120450US2

Fig. 11

12/47

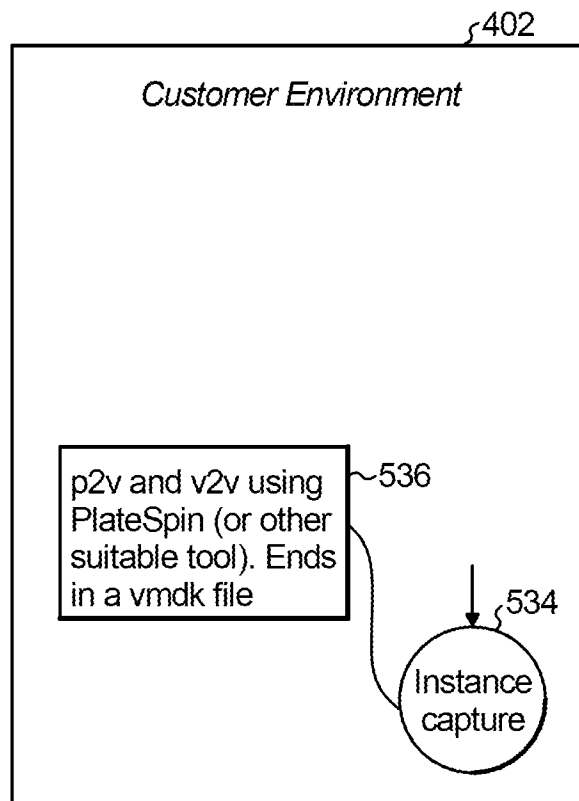


Fig. 12

13/47

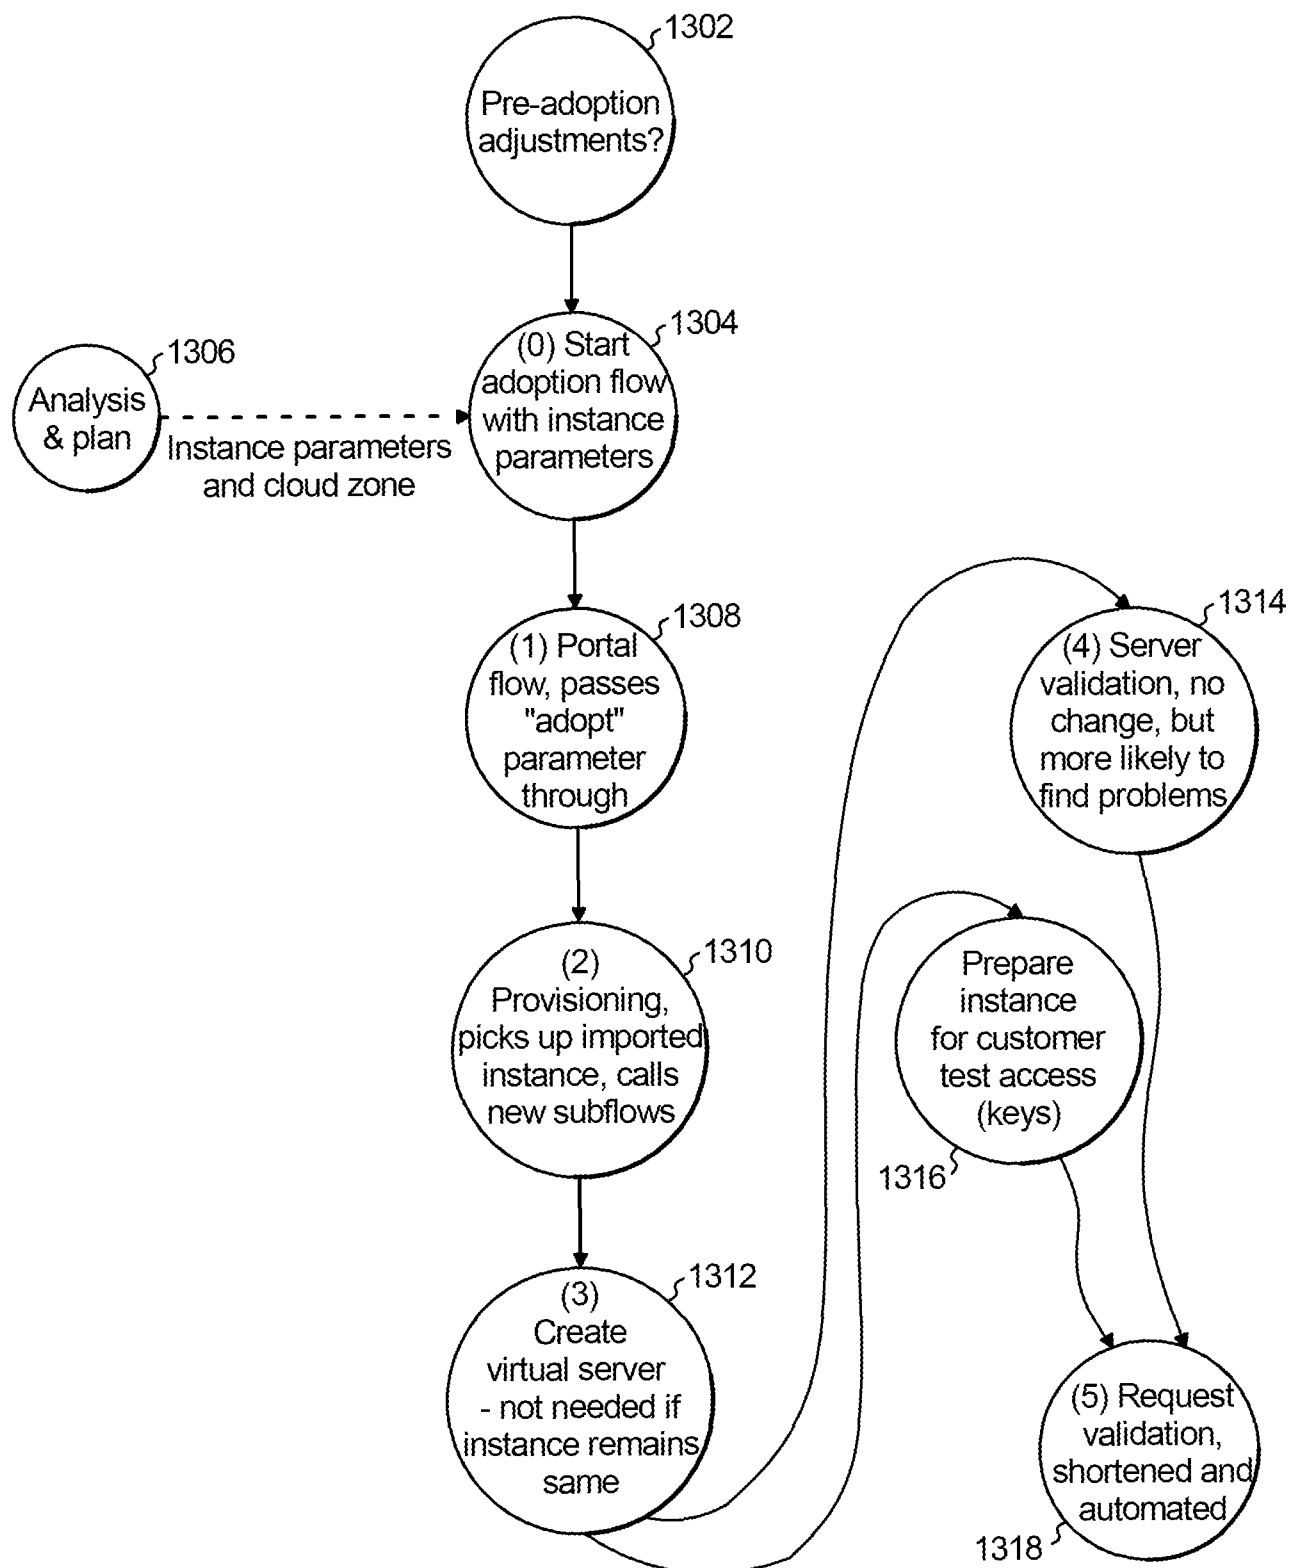


Fig. 13

1402

Self Service Center

Bulletins: (0) ▼Go To Reports Start Center Profile Sign Out Help

Home : ☰ Cart Search My News

Home Request a New Service Migration Management

Create server
Create server from migration image, not be available for creation, until the with the image was de-provisioned.

Create Server

Offering: Migrate Server

Description: Submit this request to create server from the migration image

Project:

Configuration Set (#): ▼

Set 1

Image : ☐

Location : ☐

VM Size : 2 CPUs, 8GB, 128G, System X

View My News...

Show All My Requests...

14/47

1404

Fig. 14

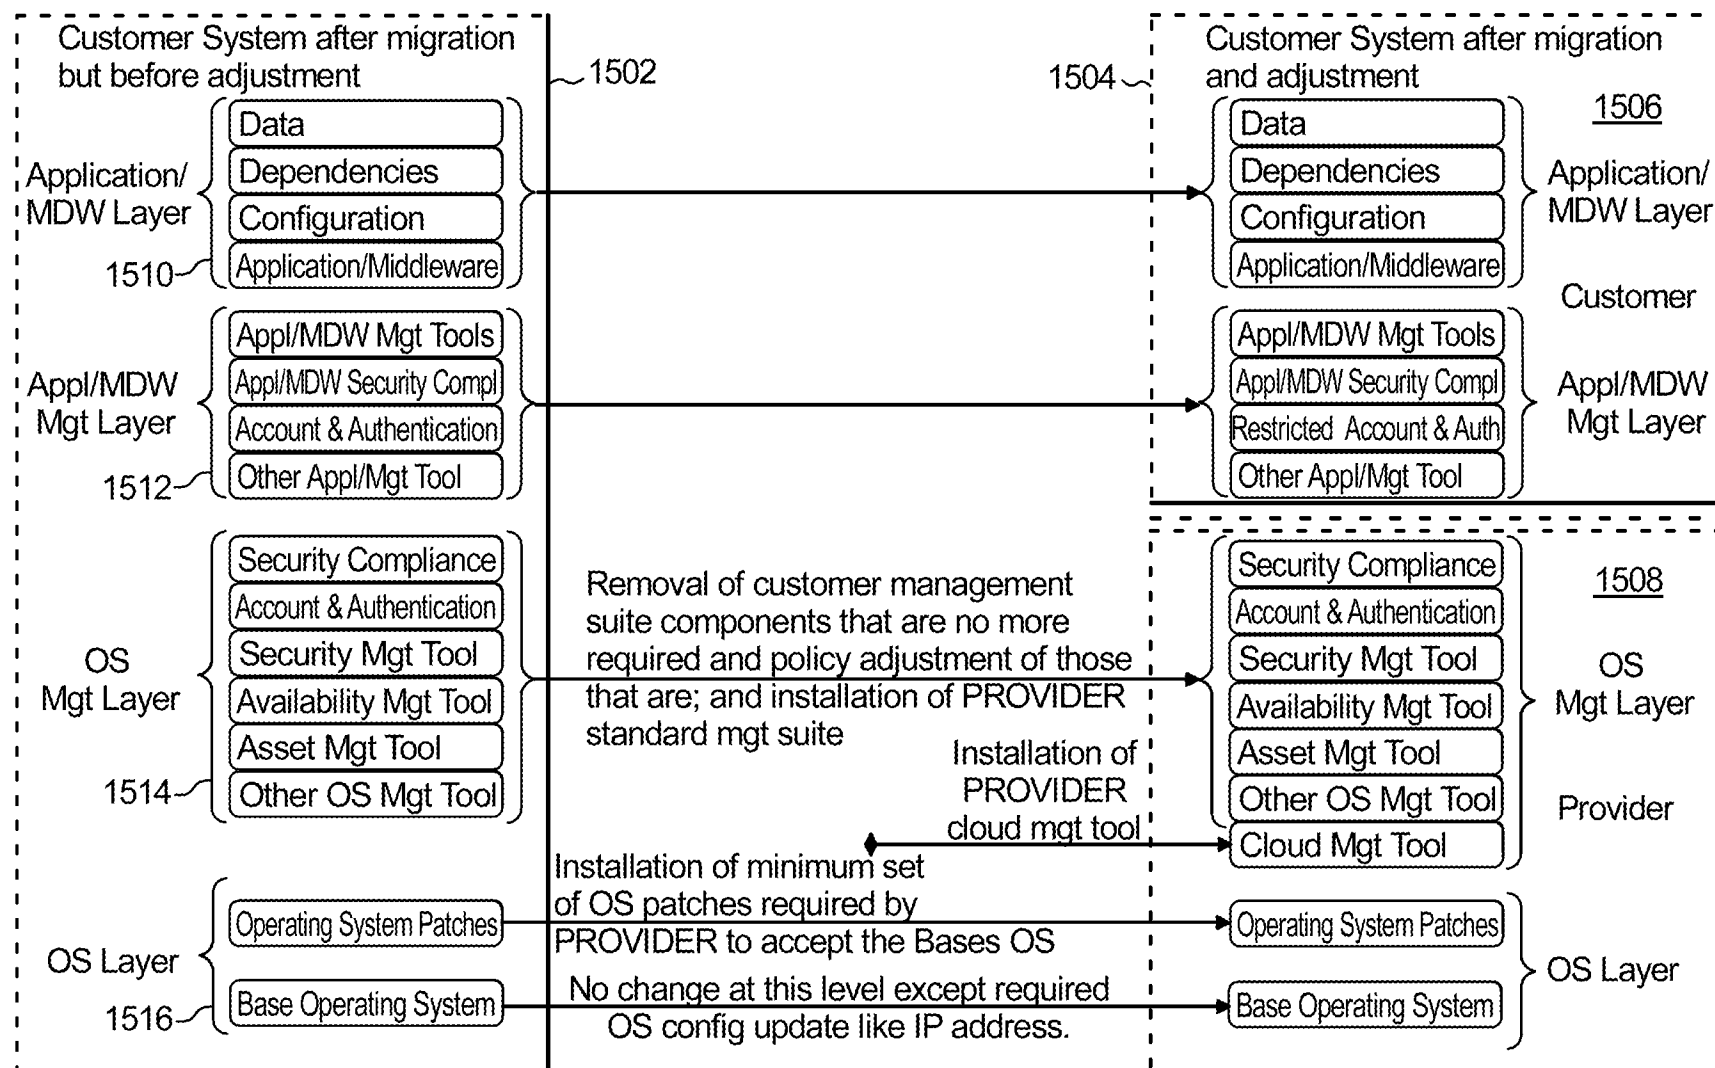


Fig. 15

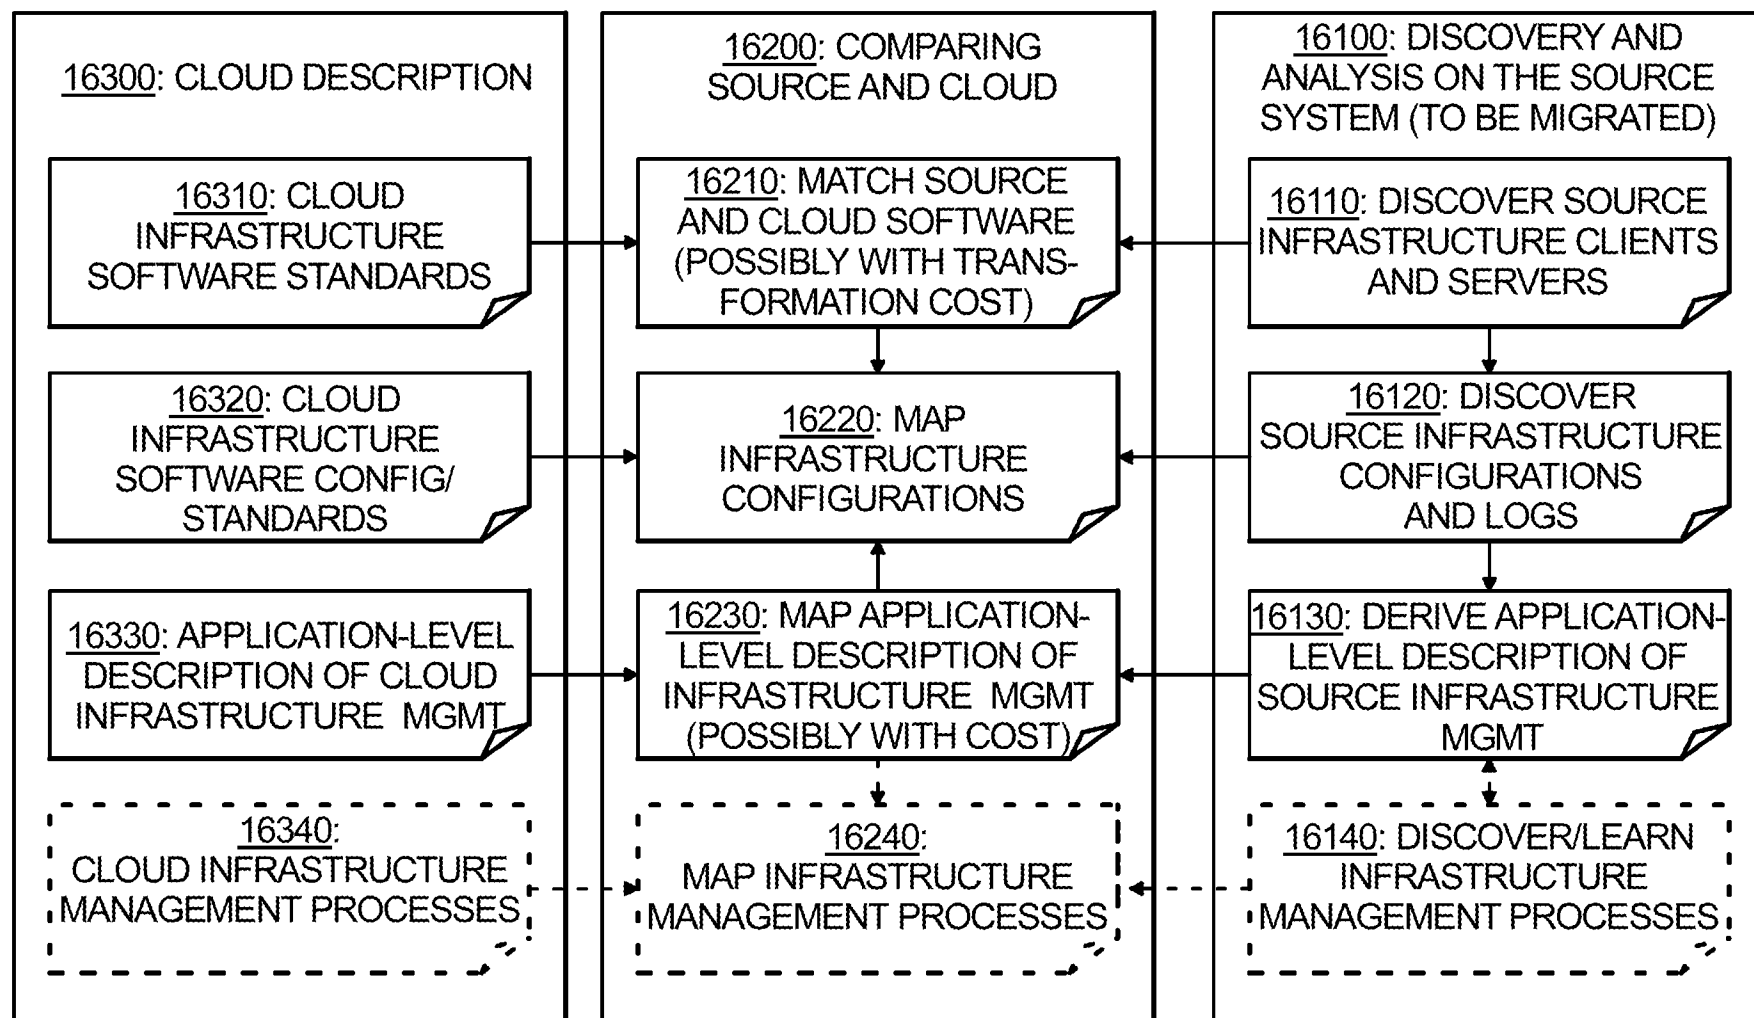


Fig. 16

16/47

YOR920120450US2

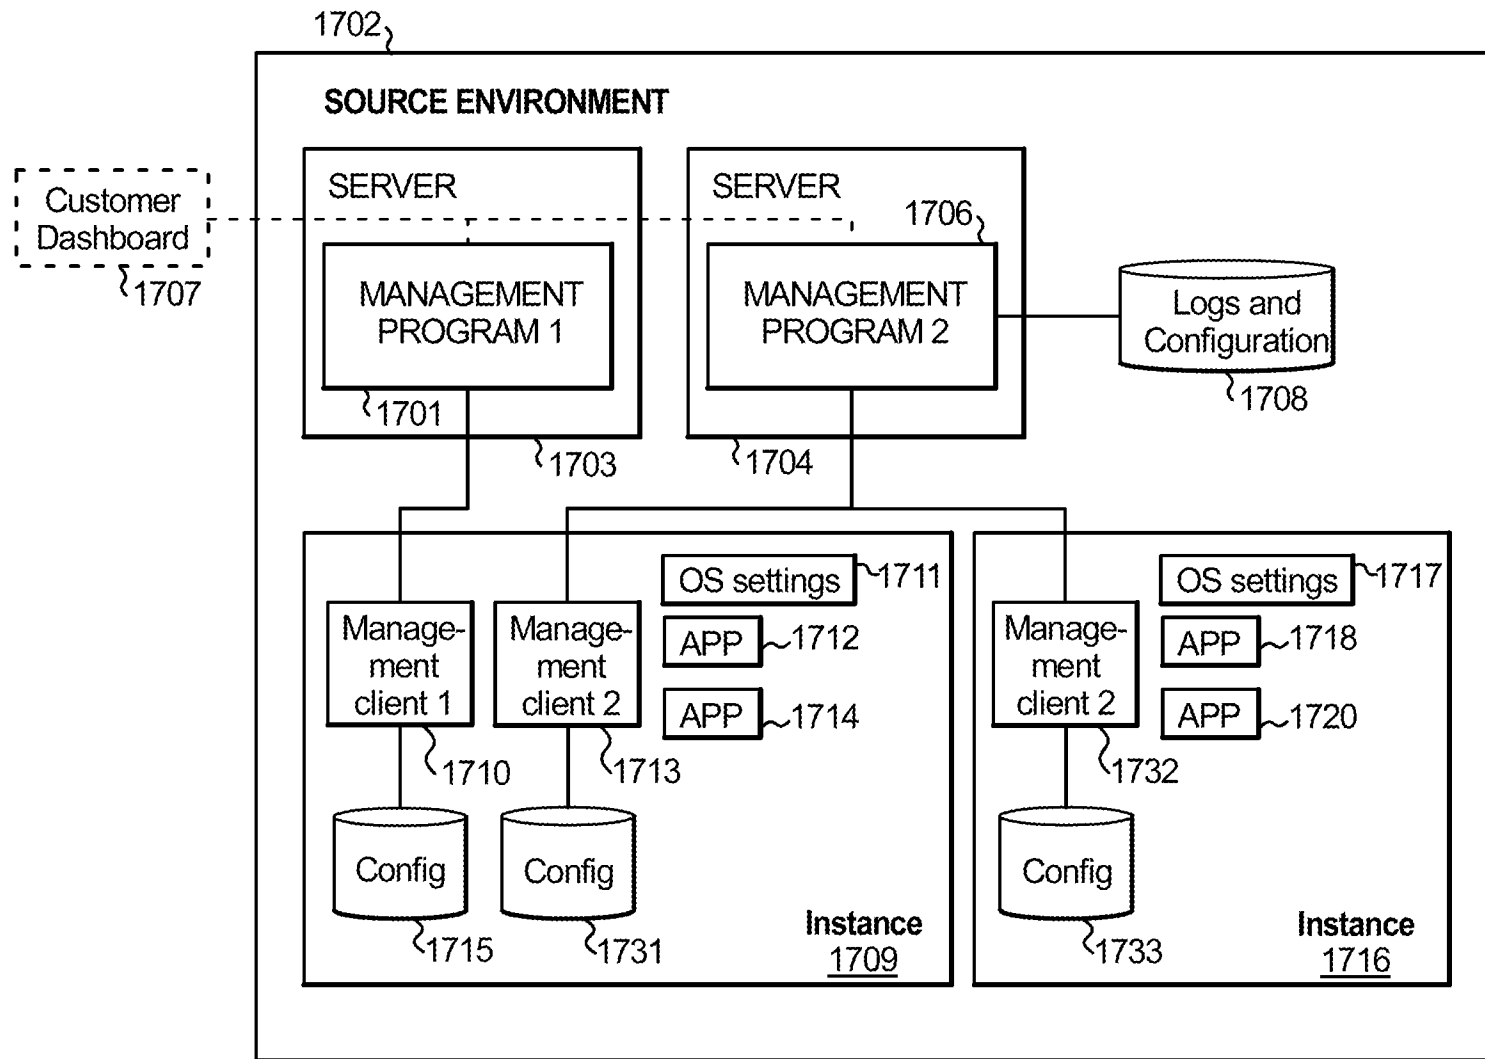


Fig. 17

17/47




Fig. 18

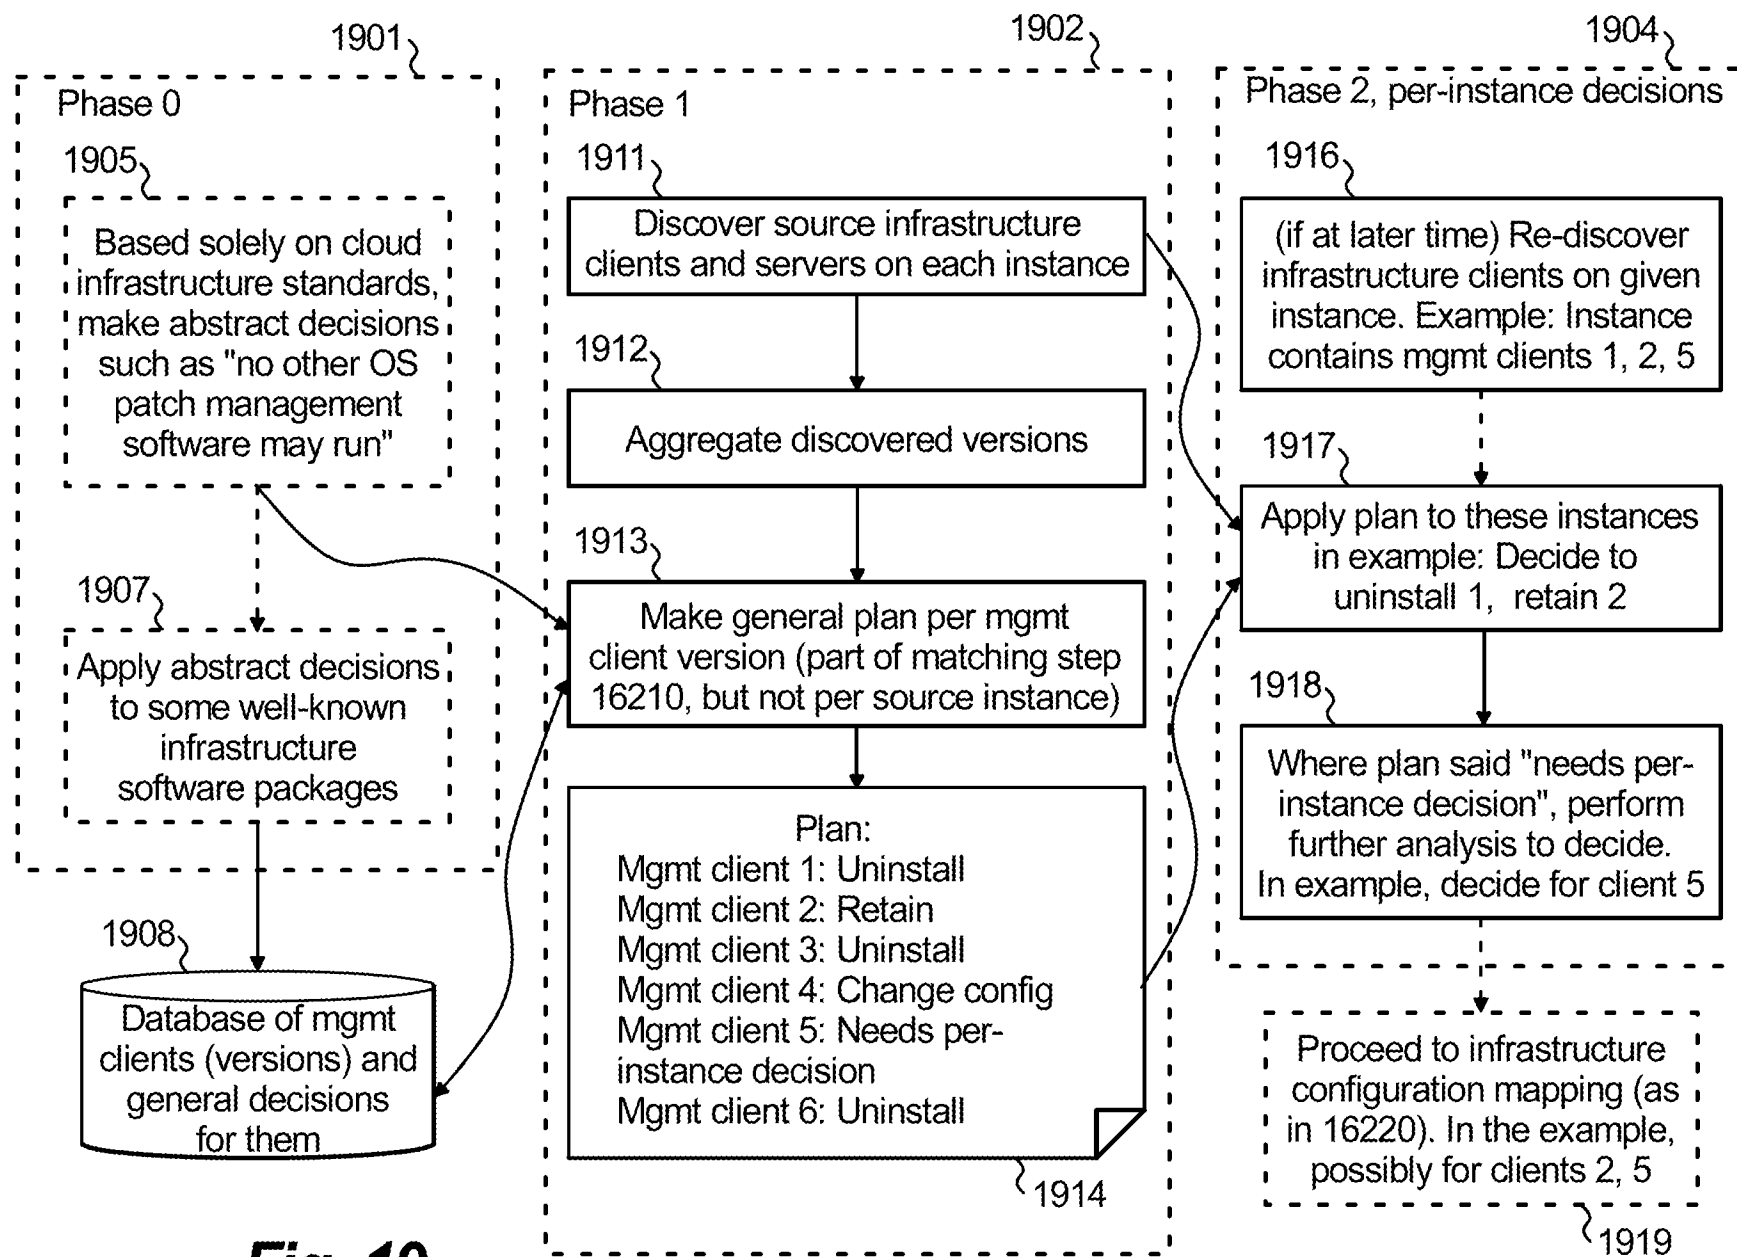


Fig. 19

19/47

YOR920120450US2

2002

2004

Management Middleware Compatibility

Select Hypervisor Name: *

IBM SCE+ ▼

Select Hypervisor Version:

1.5 ▼

Search Management Middleware compatibility information by giving

1. Middleware Name (or)
2. Middleware name and Version (or)
3. Port (or)
4. Middleware name, Version and port

Note: Please type a few characters of a middleware product name and select the middleware name from the list or Select version from the dropdown list or Select port from the list and click on Search button.

Middleware Name:

HP Asset Manager ▼

Middleware Version:

▼

Port:

▼

Search

Management Middleware Compatibility

Select Hypervisor Name: *

IBM SCE+ ▼

Select Hypervisor Version:

1.5 ▼

Search Management Middleware compatibility information by giving

1. Middleware Name (or)
2. Middleware name and Version (or)
3. Port (or)
4. Middleware name, Version and port

Note: Please type a few characters of a middleware product name and select the middleware name from the list or Select version from the dropdown list or Select port from the list and click on Search button.

Search Results:

HP Asset Manager is Functional Conflict with IBM SCE+ 1.5 because the hypervisor runs functionality equivalent Tivoli Asset Discovery for Distributed 7.7.2 FP1

Search again:

Middleware Name:

▼

Middleware Version:

▼

Port:

▼

Search

20/47

YOR920120450US2

Fig. 20

Command Editor 1 - DB2COPY2

Command Editor Selected Edit View Tools Help

Commands Query Results Access Plan

Edits to these results are performed as positioned UPDATES and DELETES. Use the Tools Settings notebook to change the form of editing.

HOST_NAME	MW_VENDOR	MW_DISTRIBUTION_NAME	FUNCTION
NMLECDBM	IBM	Tivoli Monitoring UnixLog Agent	UNINSTALL
NMLECDBP	IBM	Tivoli Monitoring UnixLog Agent	UNINSTALL
NMLECDBQ	IBM	Tivoli Monitoring UnixLog Agent	UNINSTALL
NMLECDBS	IBM	Tivoli License Compliance Man...	UNINSTALL
NMLECDBT	IBM	Tivoli Security Compliance Man...	UNINSTALL

Add Row

Delete Row

Commit Roll Back Fetch More Rows

☐ Automatically commit updates

5 row(s) in memory

21/47

YOR920120450US2

Fig. 21





Fig. 22

23/47

**Fig. 23**

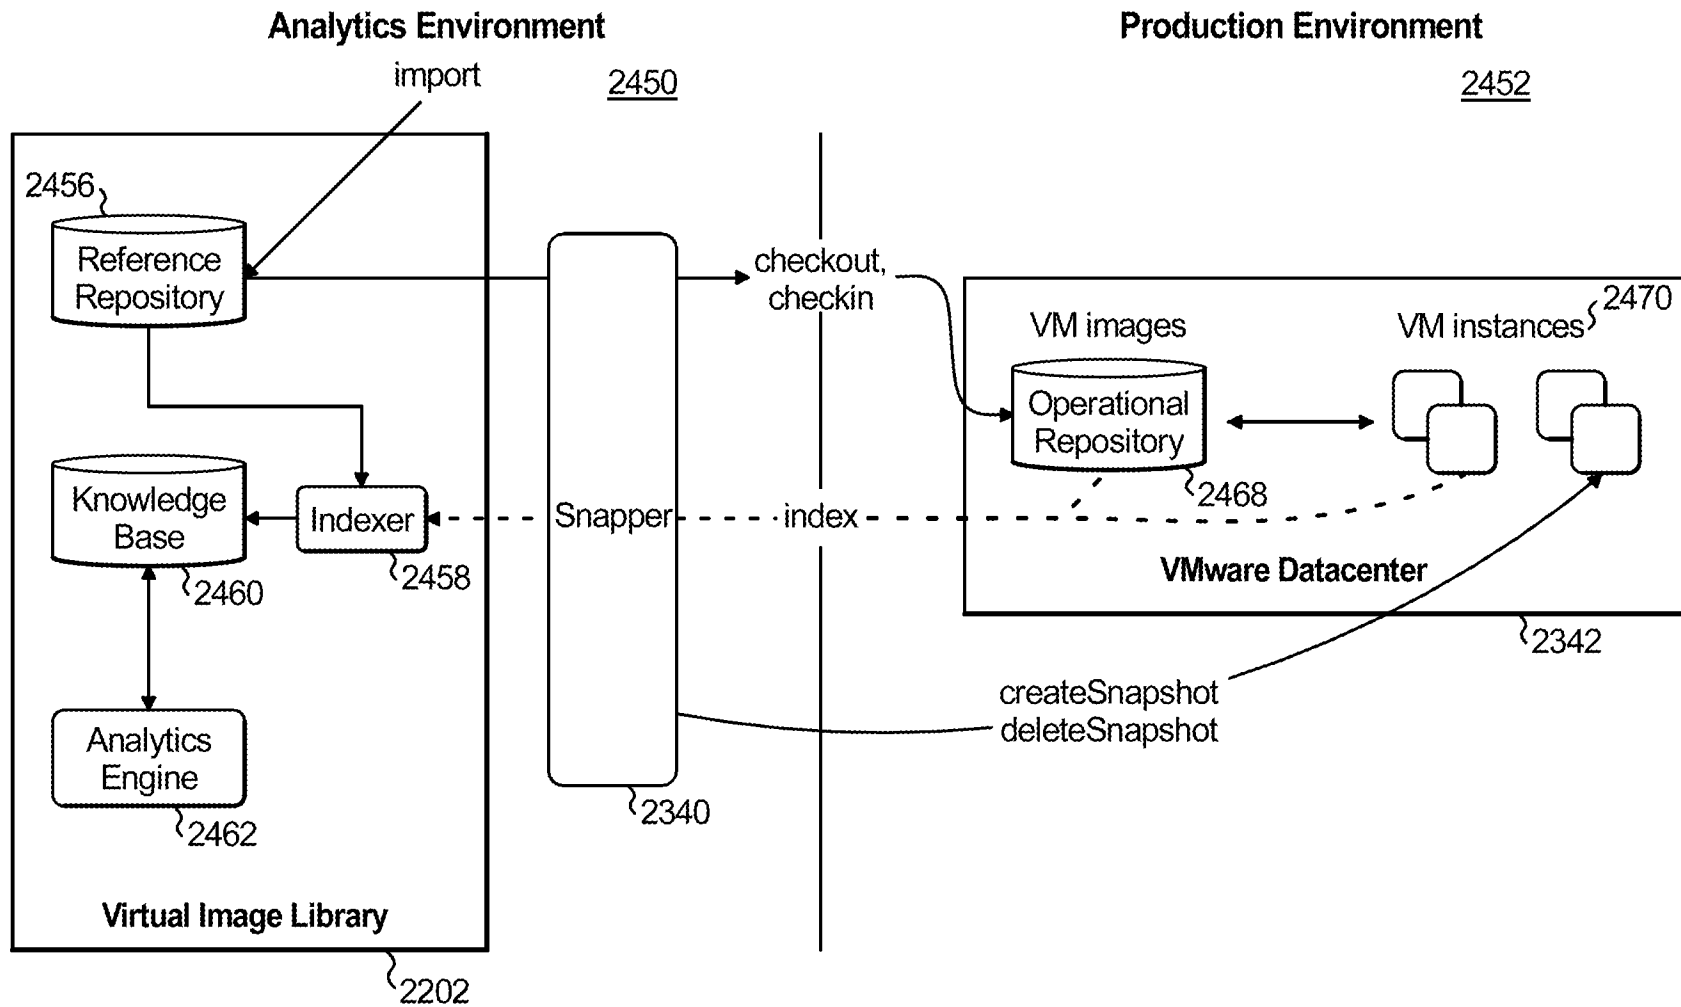

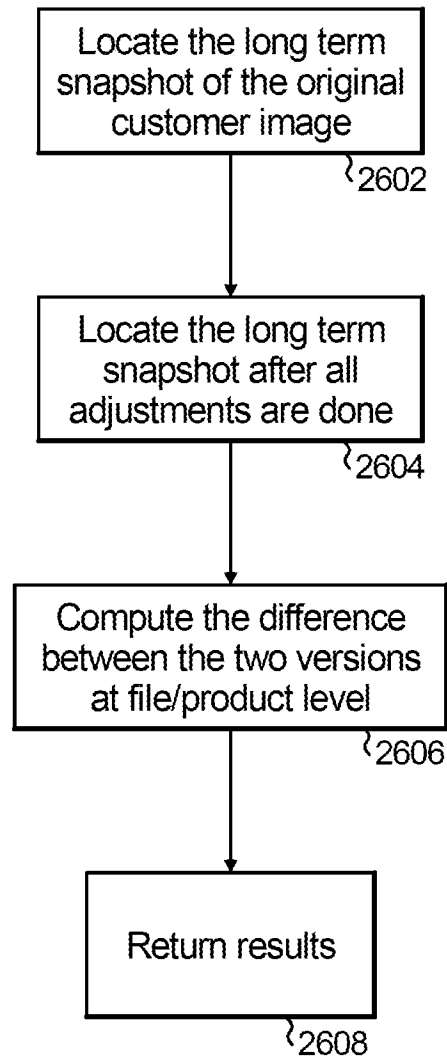
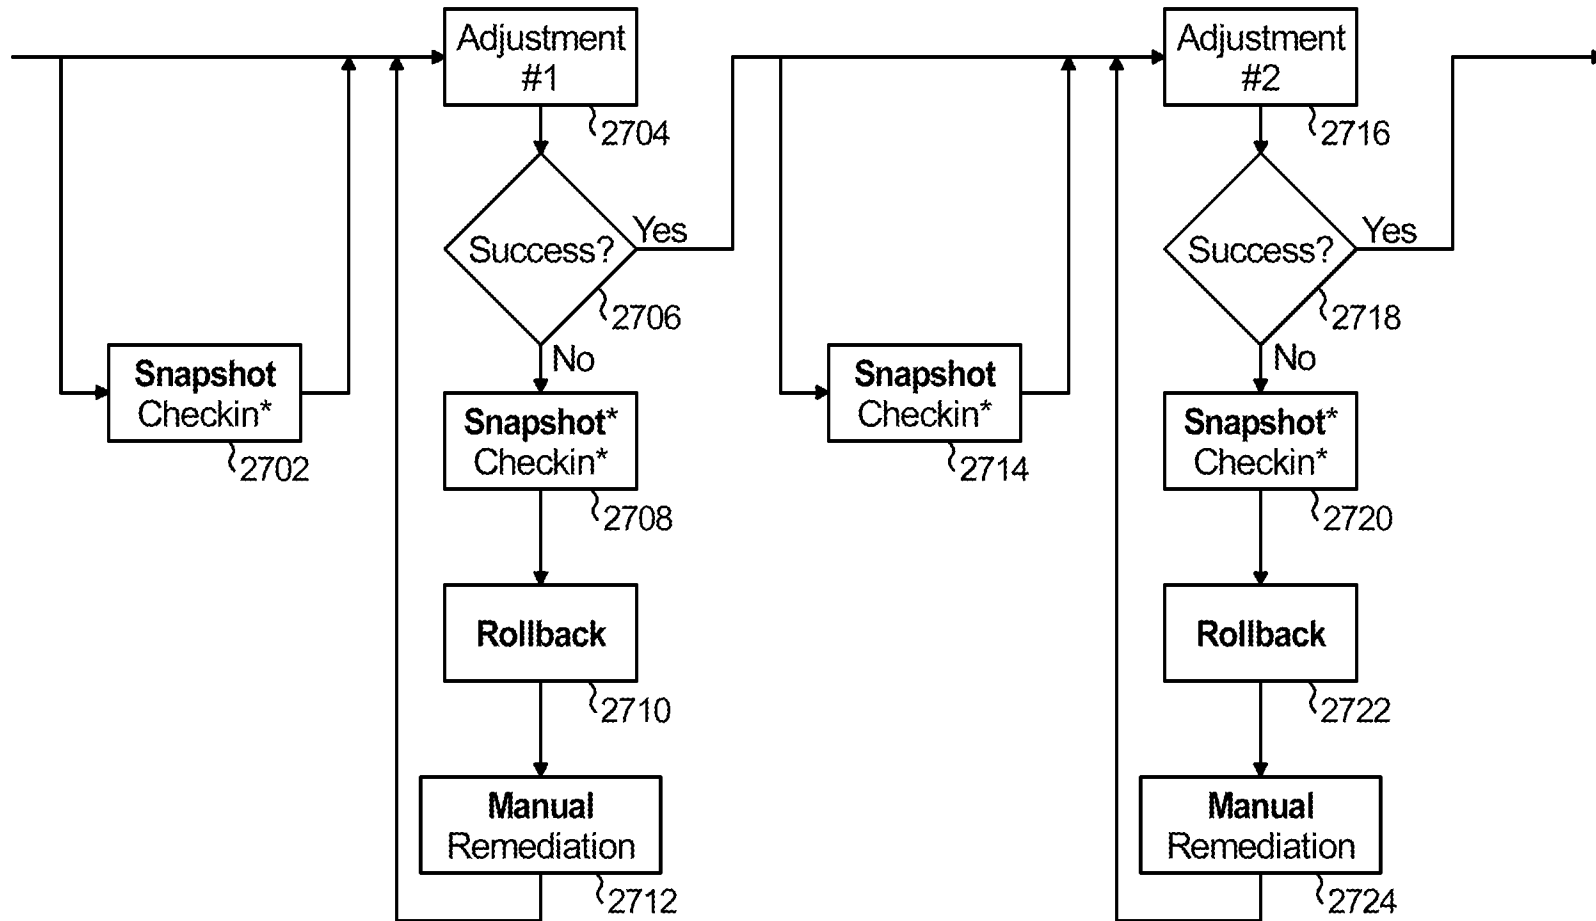


Fig. 24

25/47**Fig. 25**

26/47***Fig. 26***



Checkin*: Selective checkin
Snapshot*: Selective snapshot

Fig. 27

27/47

YOR920120450US2

28/47

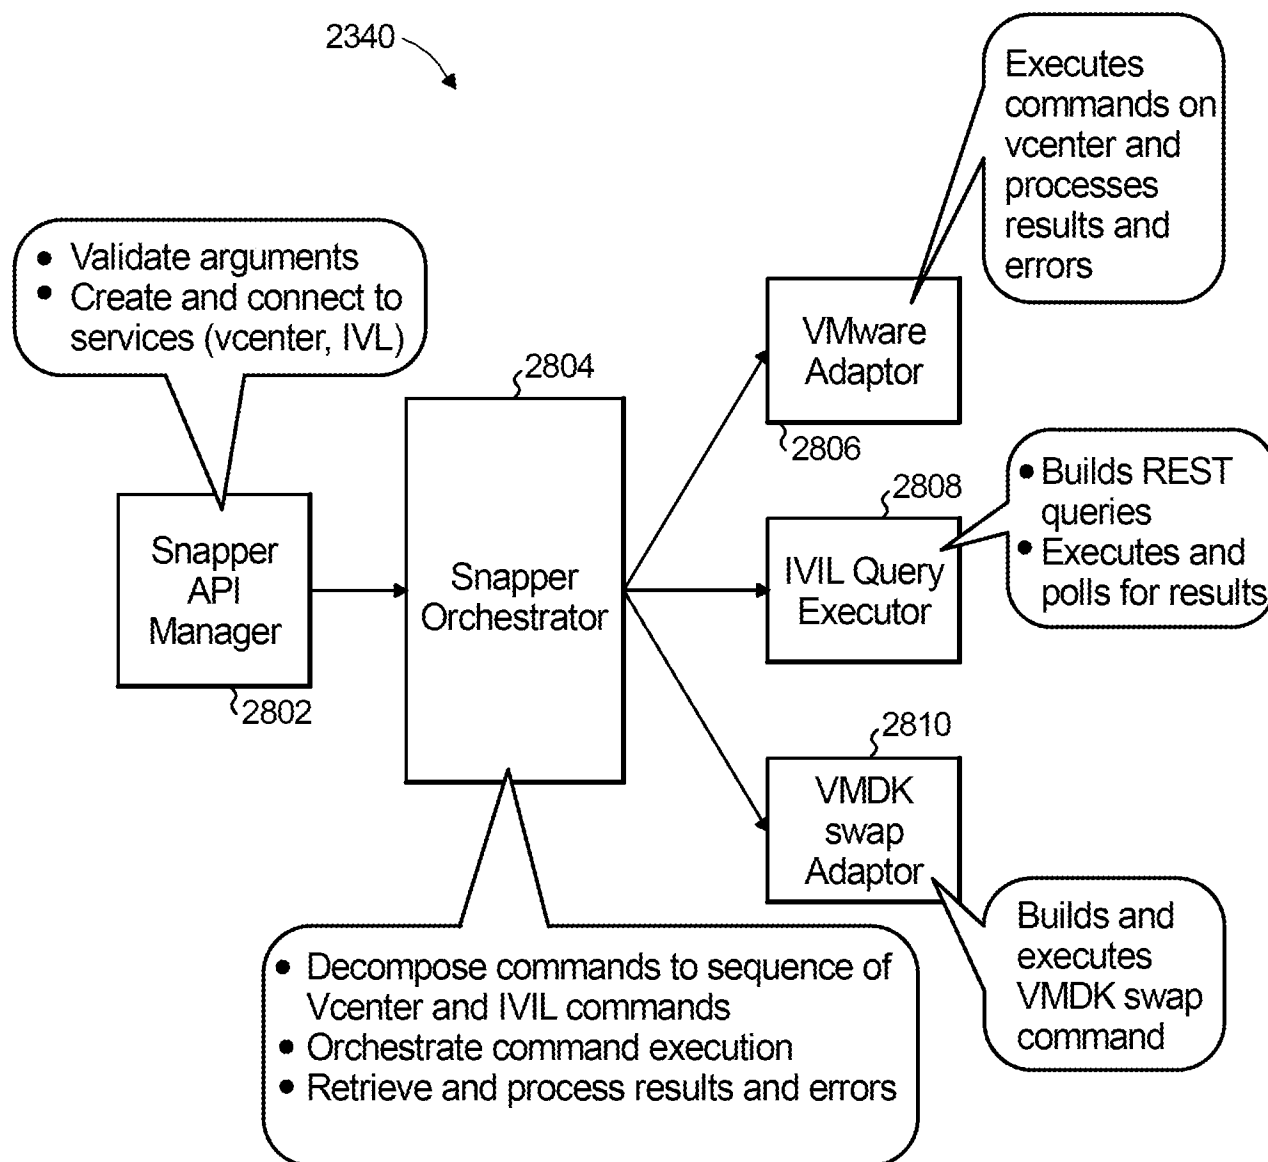


Fig. 28

29/47

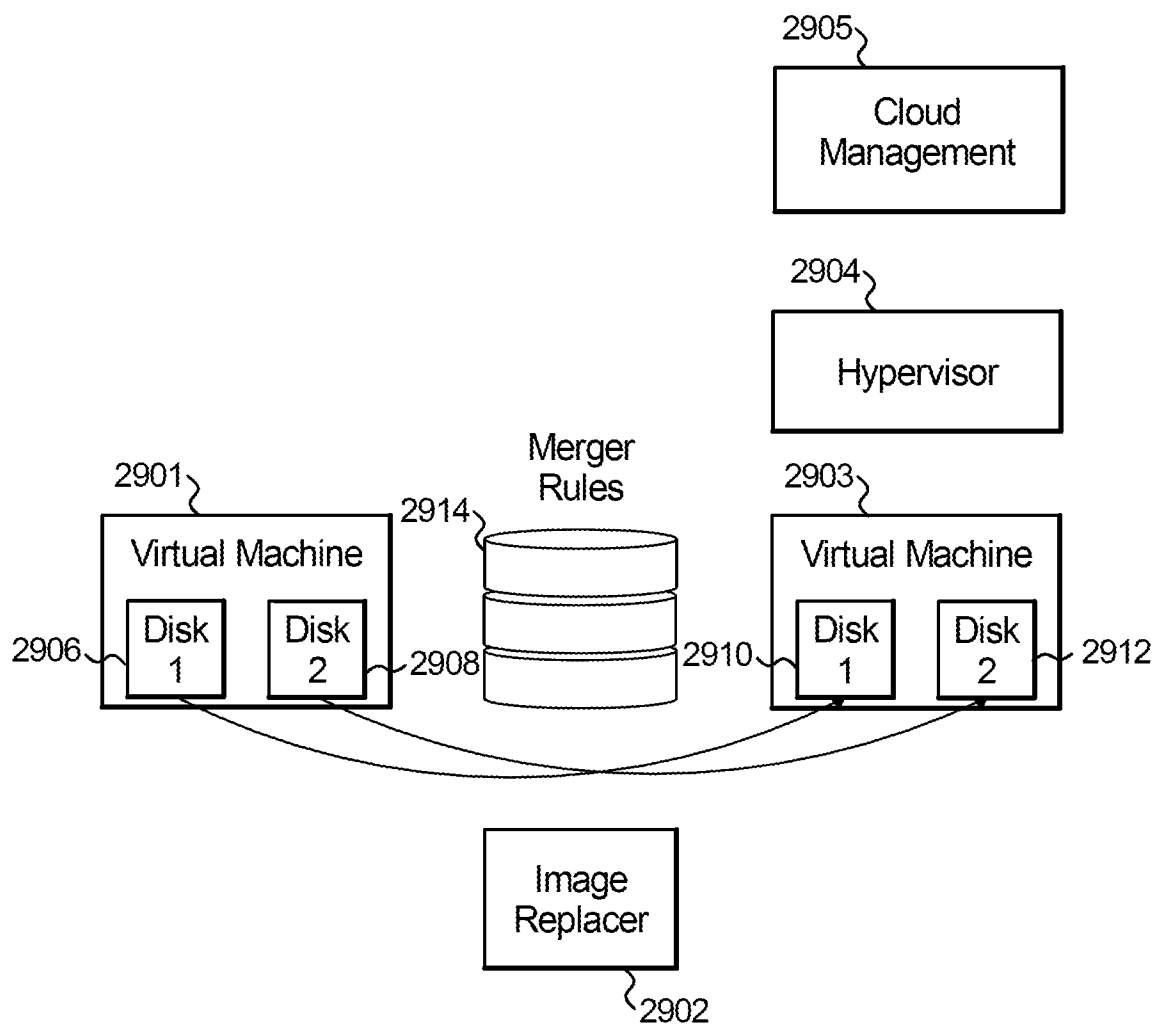


Fig. 29


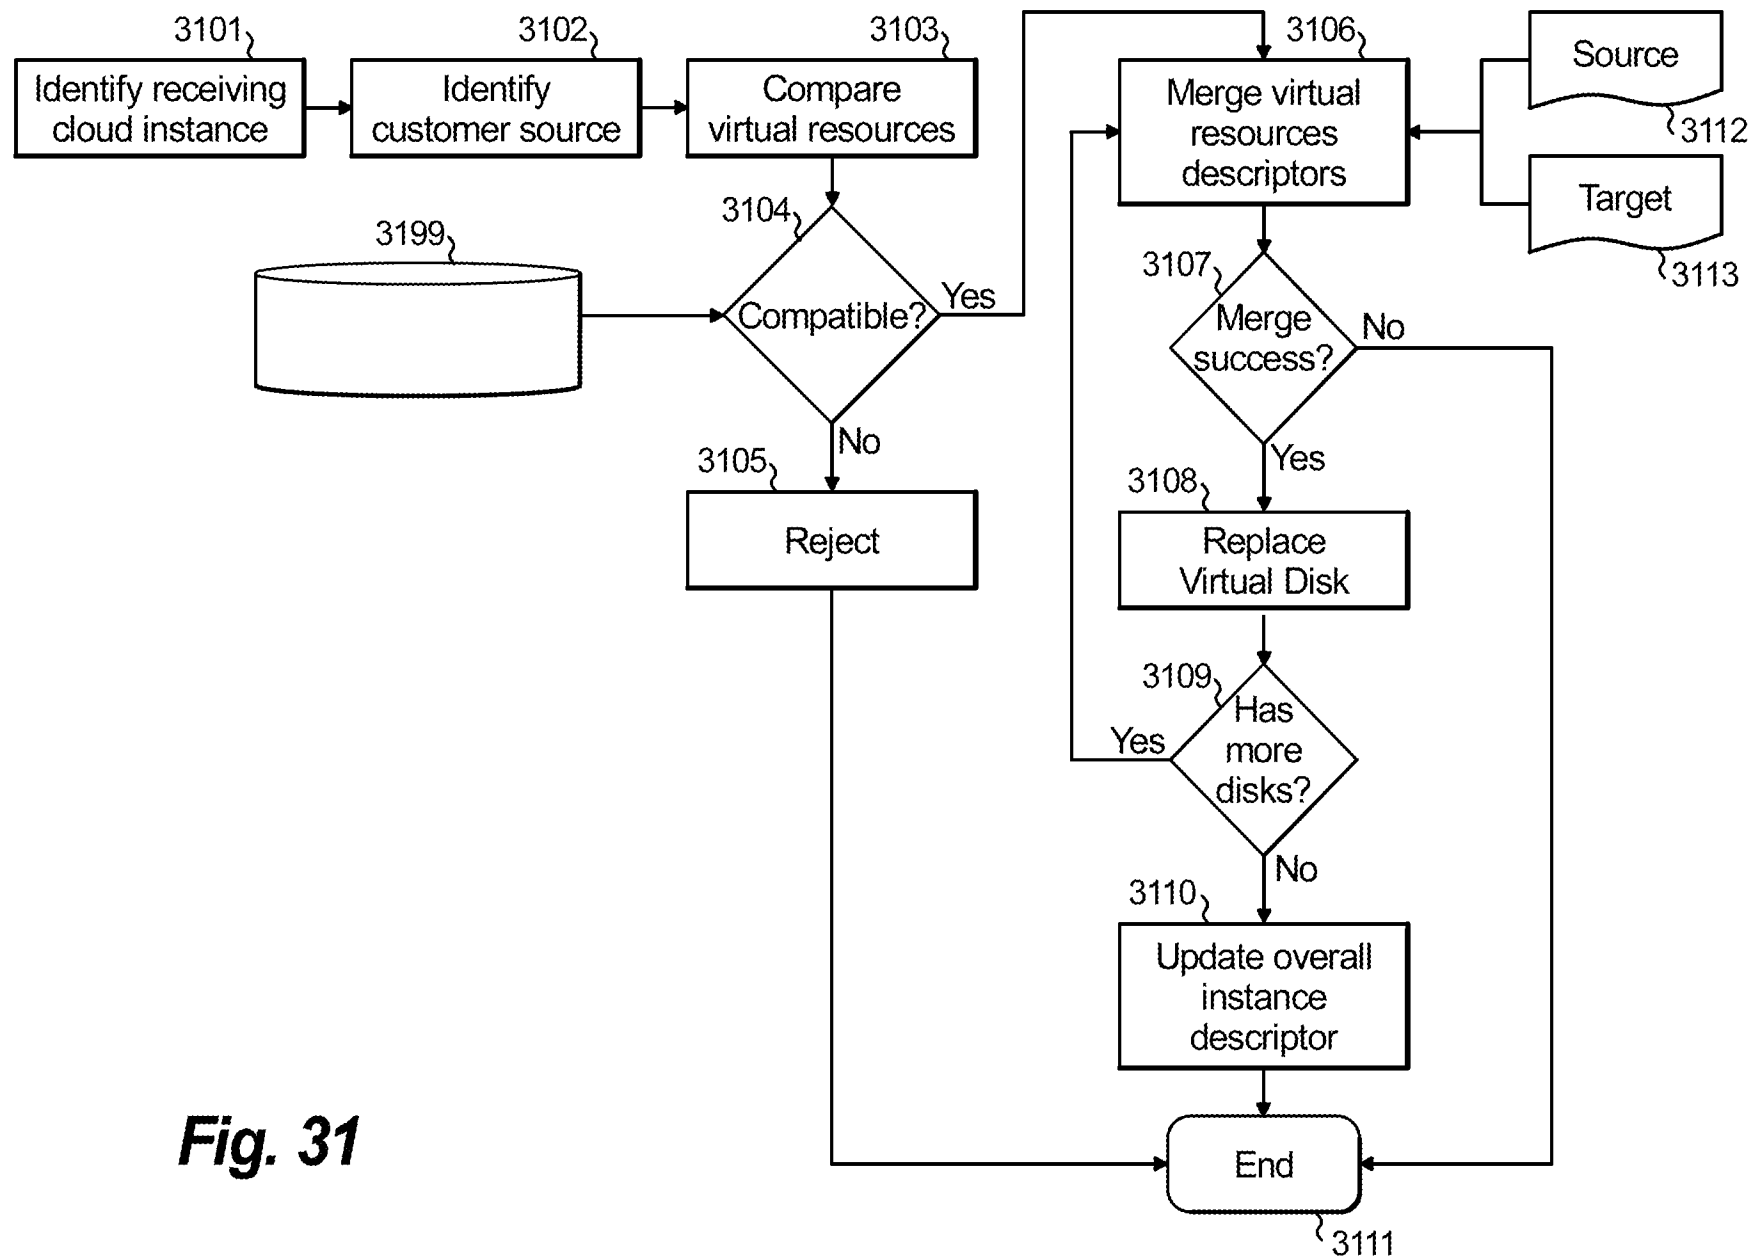


Fig. 30



32/47

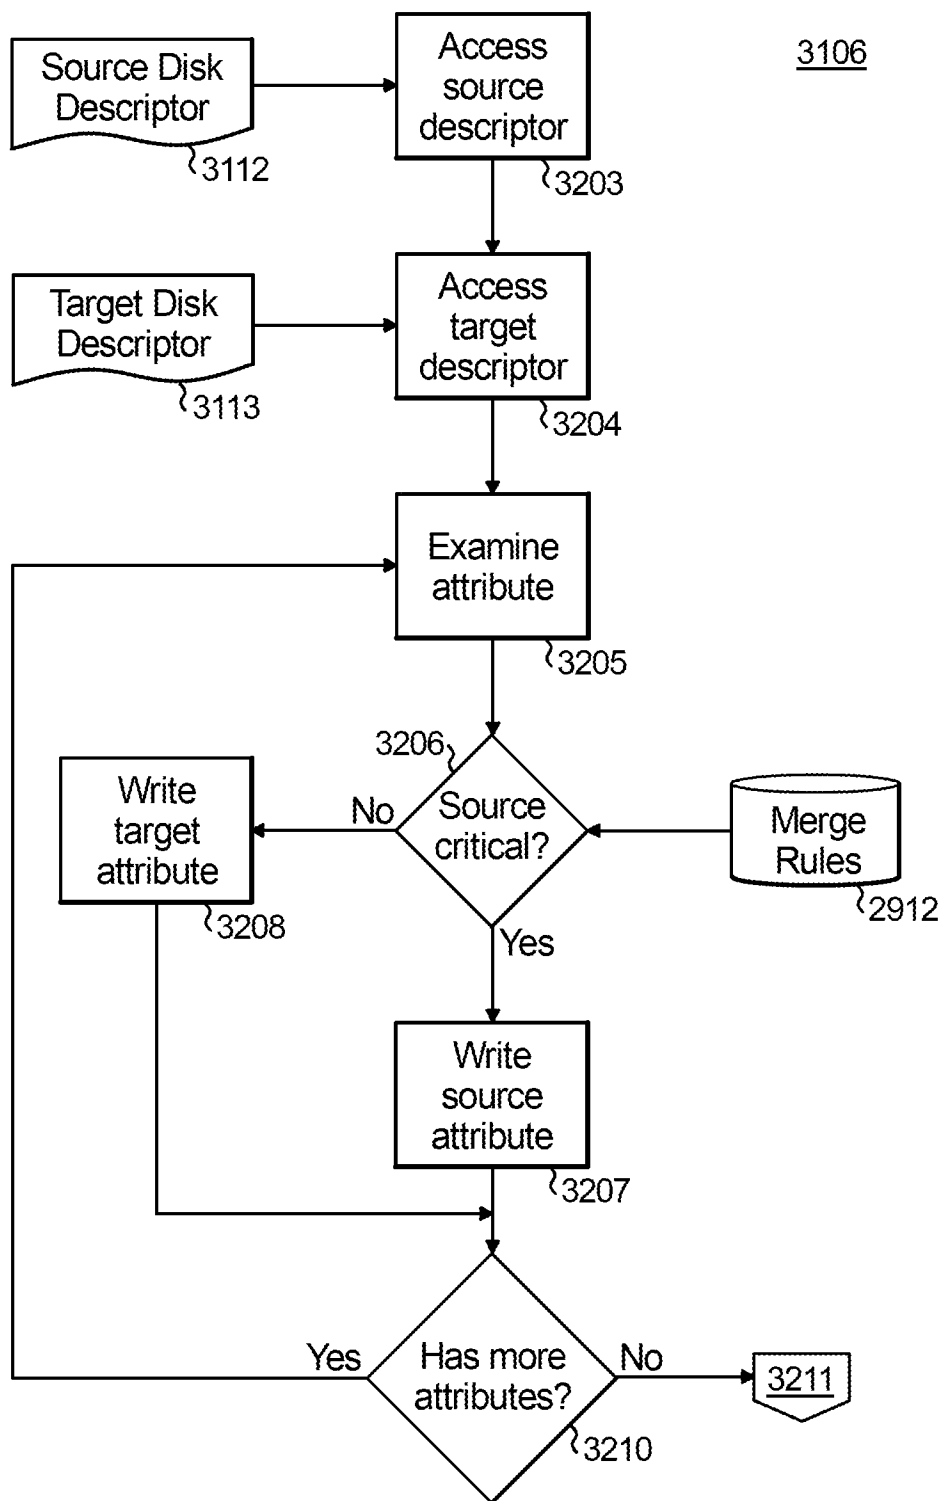


Fig. 32

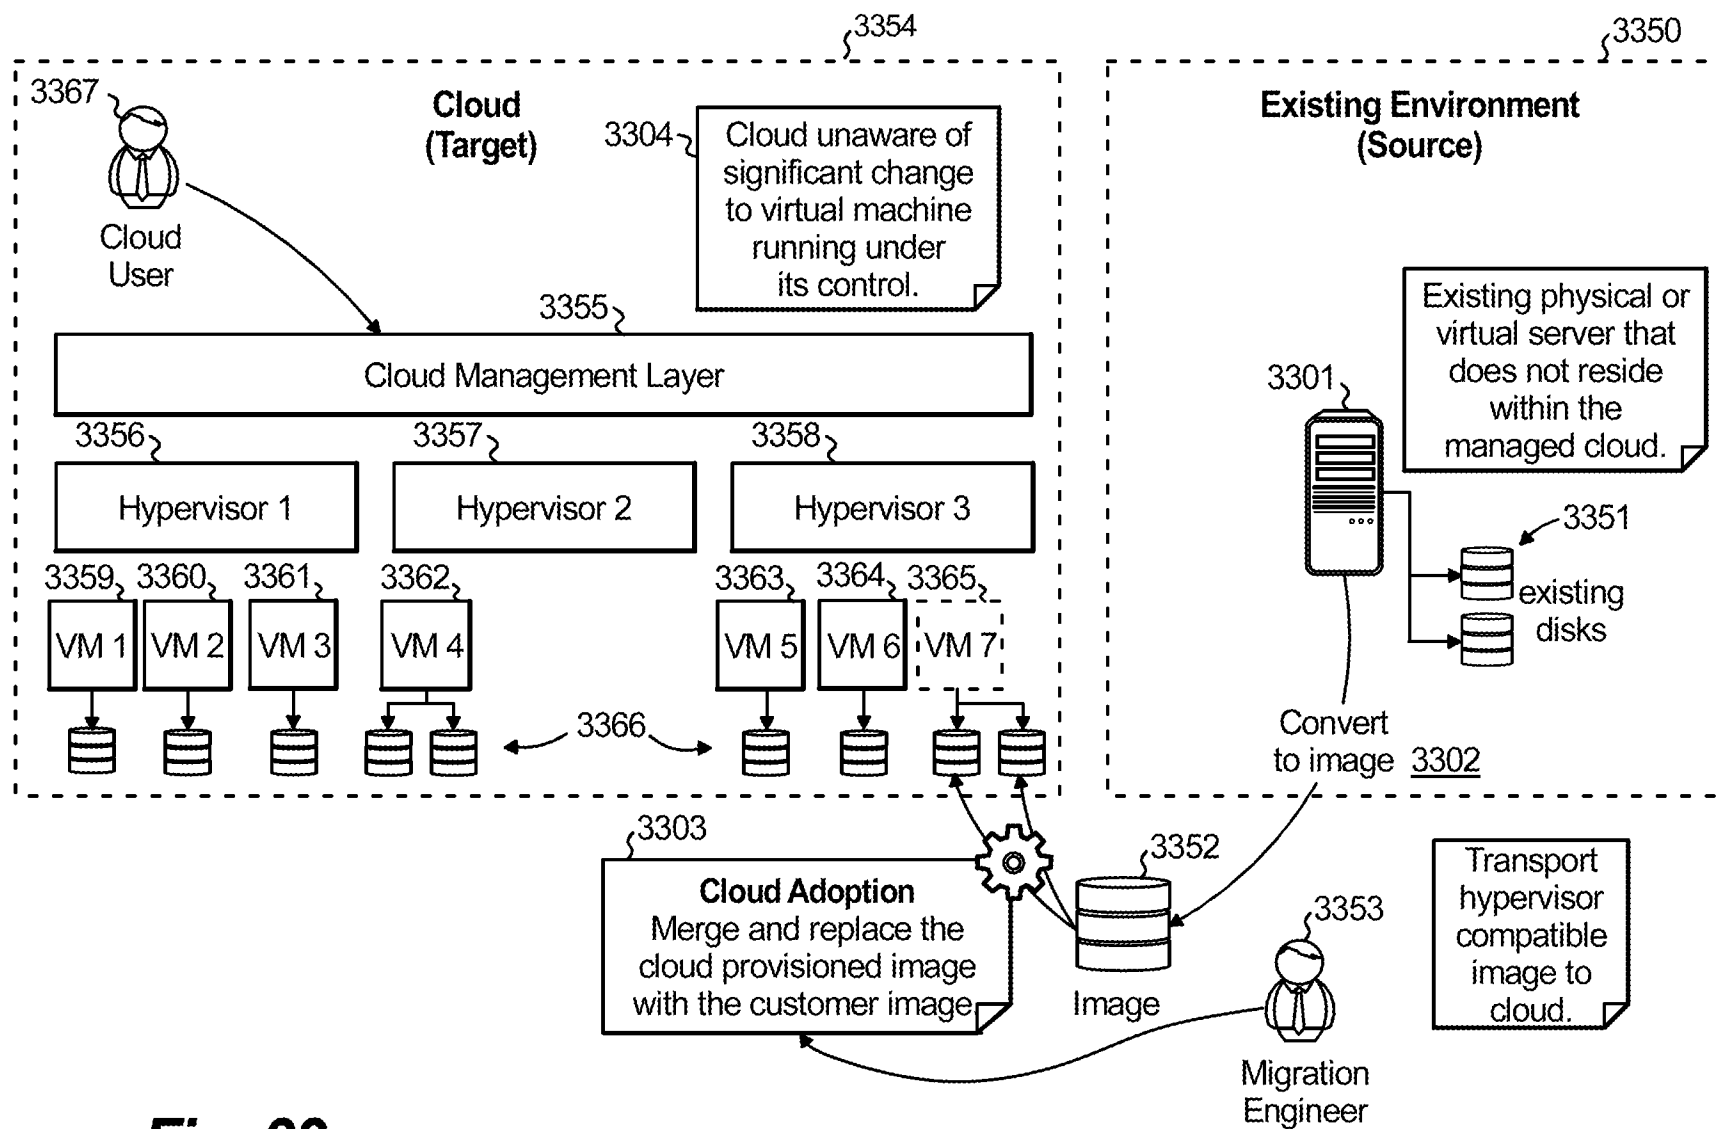
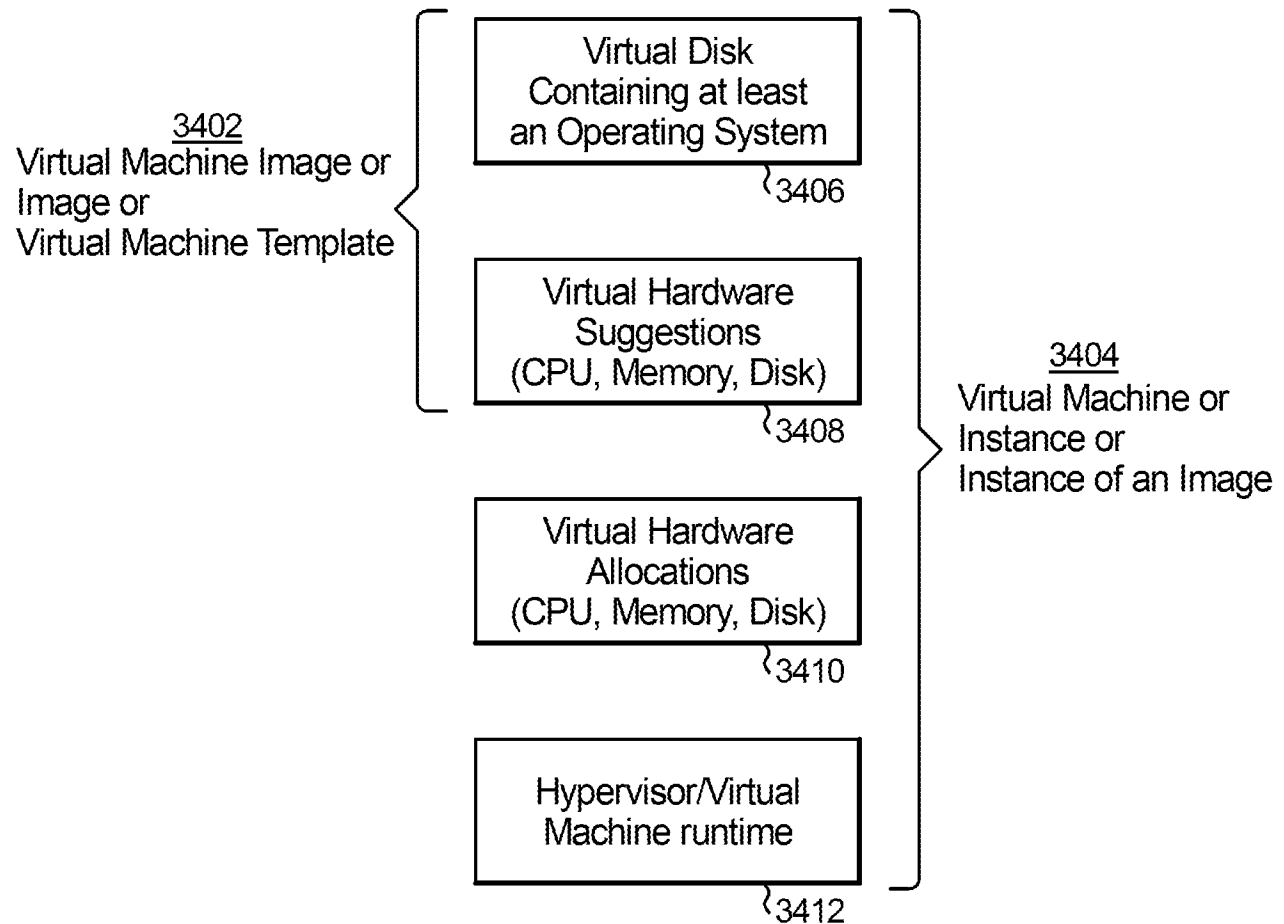



Fig. 33

34/47**Fig. 34**

35/47

**Fig. 35**

36/47

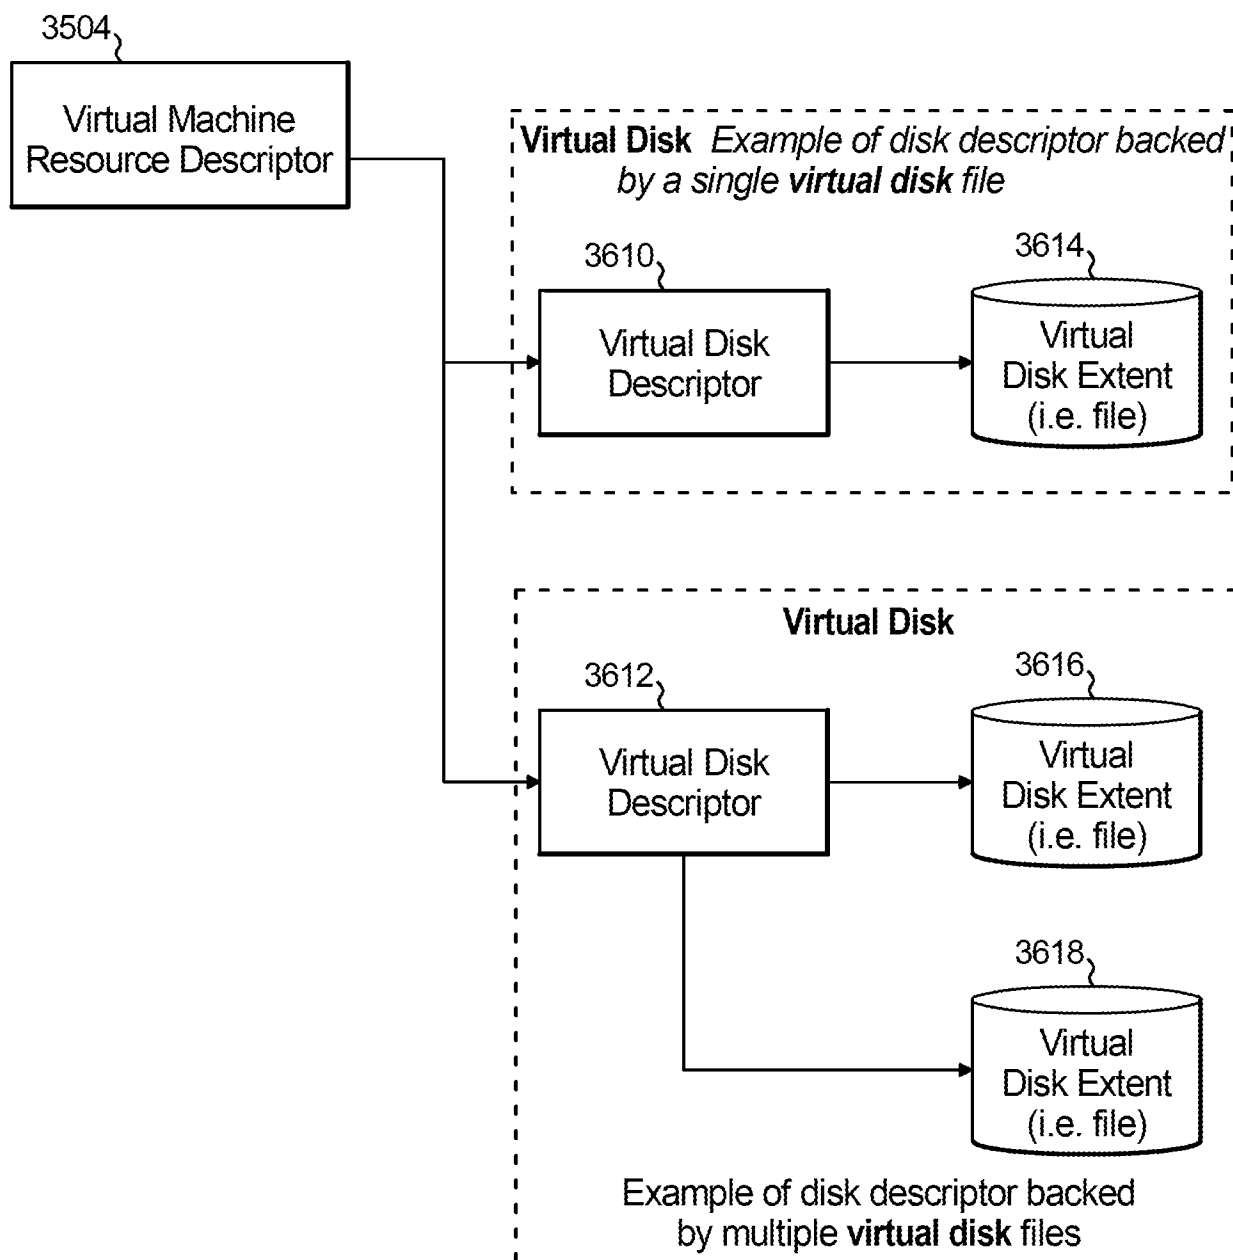
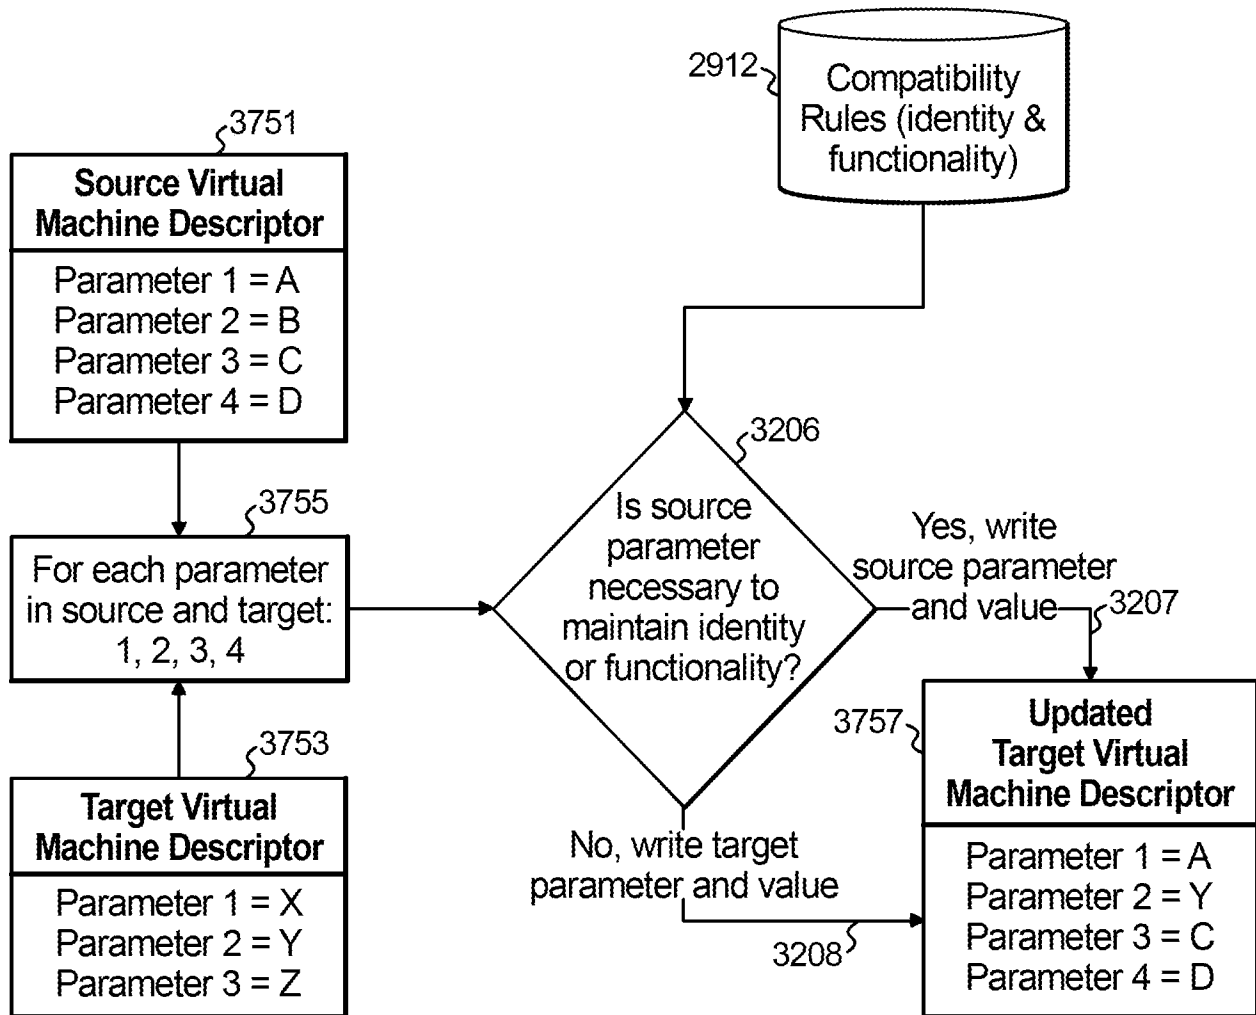
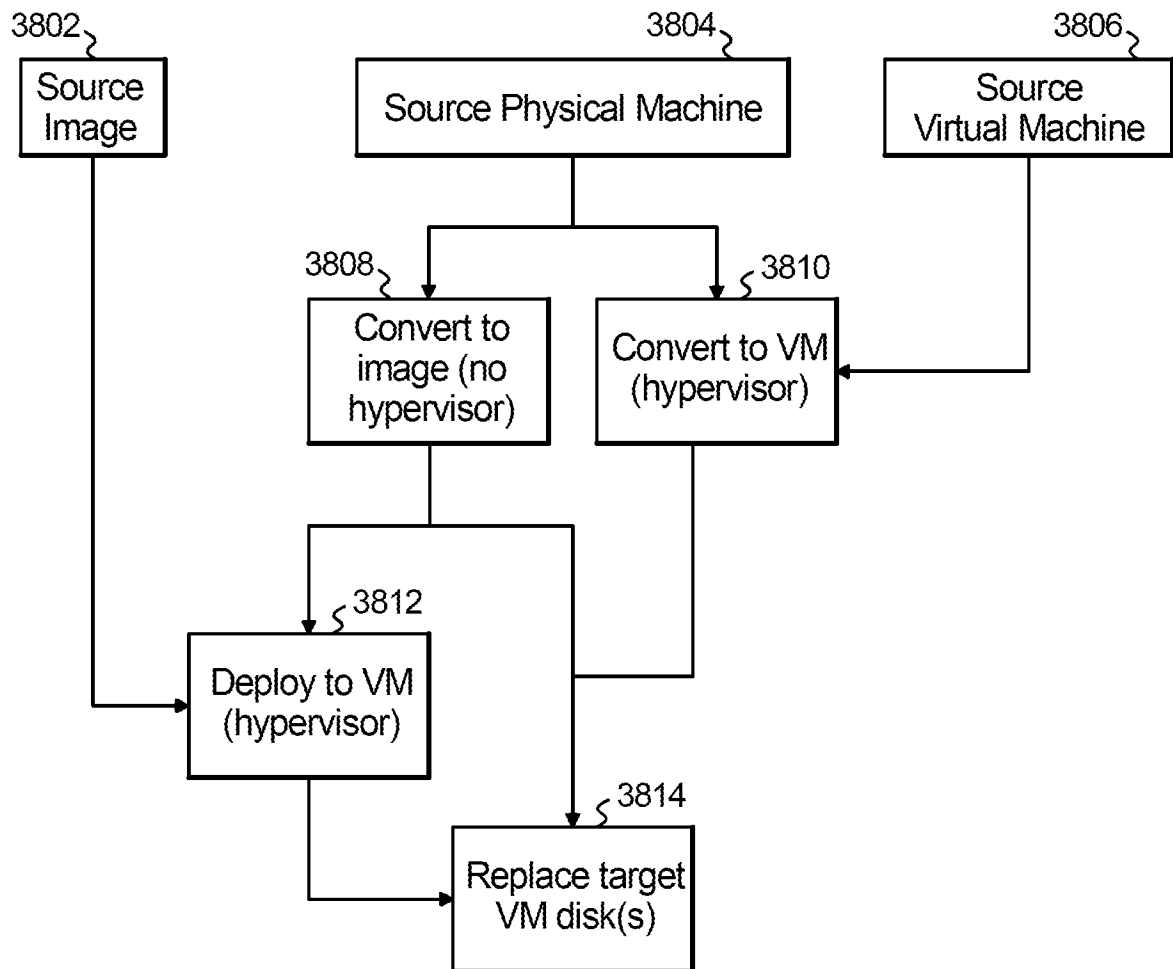


Fig. 36

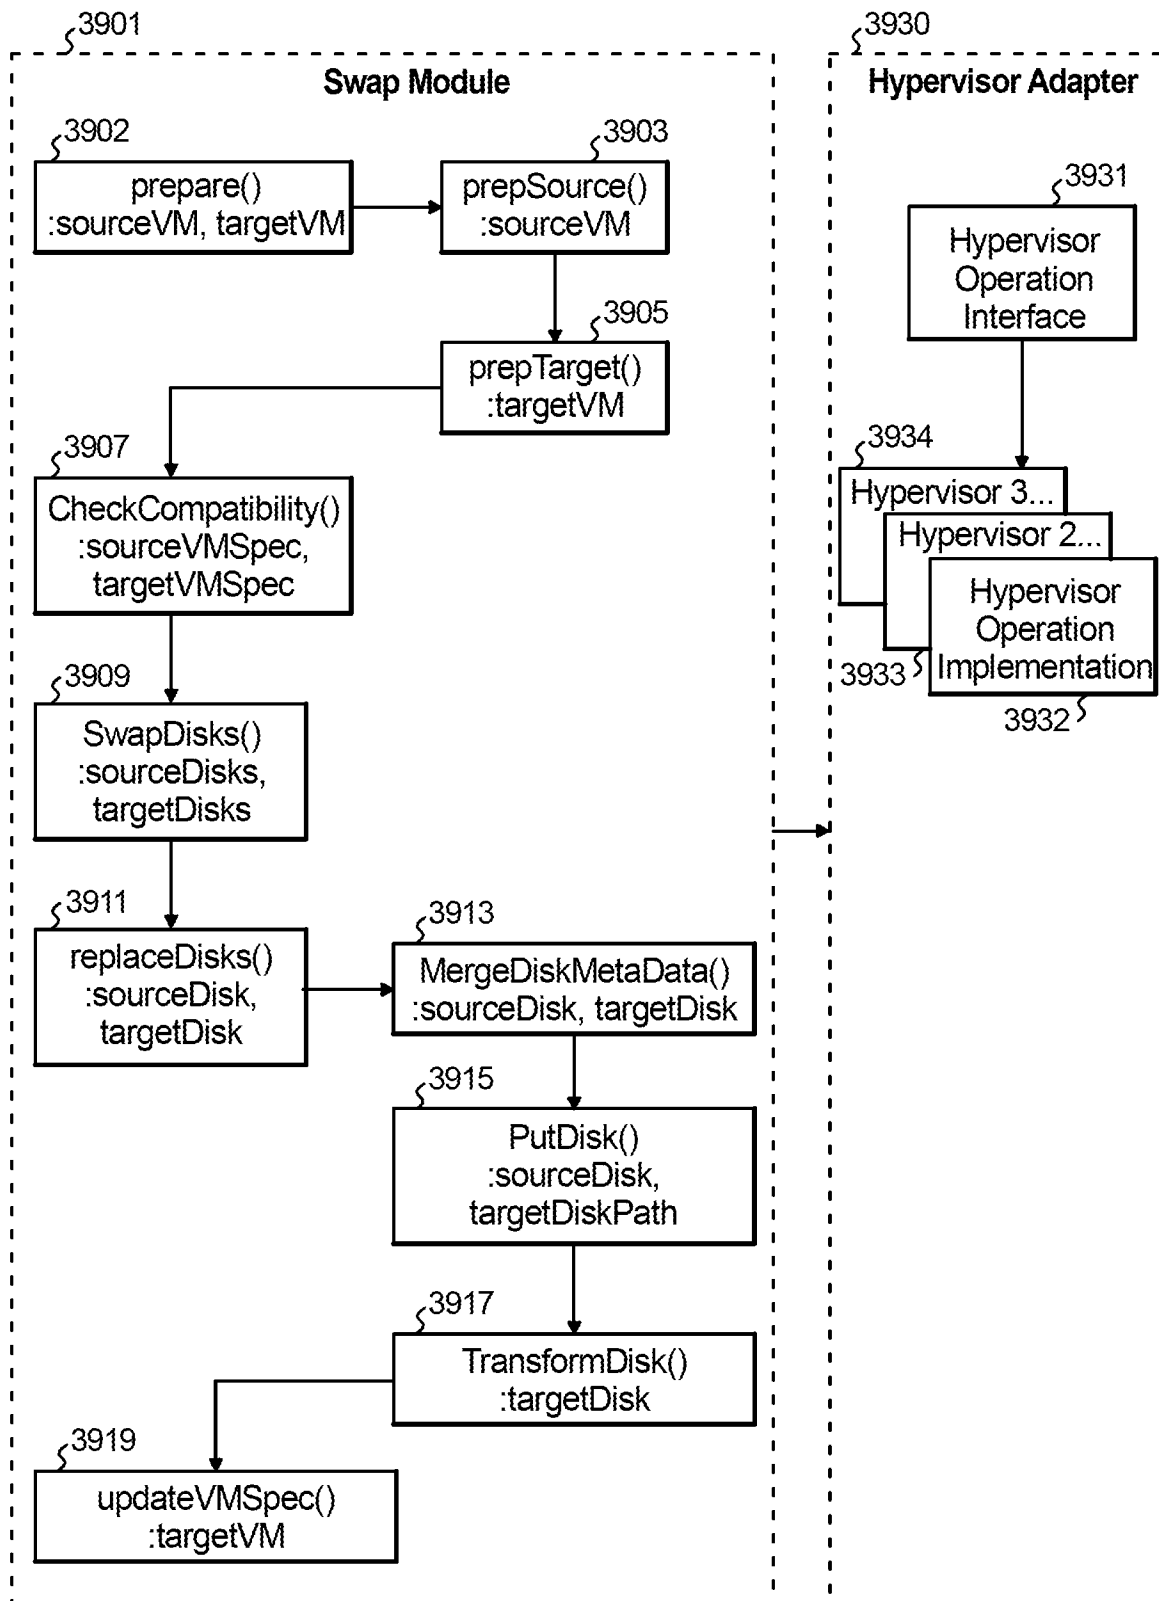
37/47

**Fig. 37**

38/47

**Fig. 38**

39/47

**Fig. 39**




Fig. 40

41/47

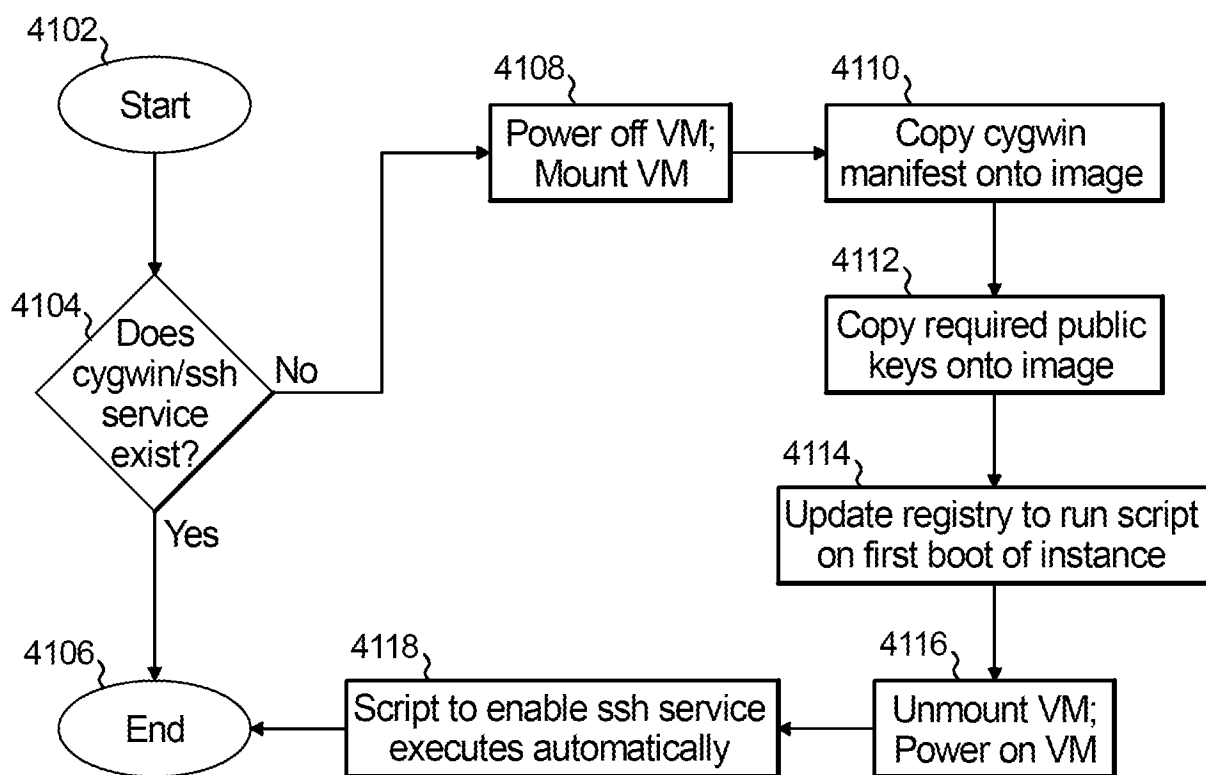


Fig. 41

42/47

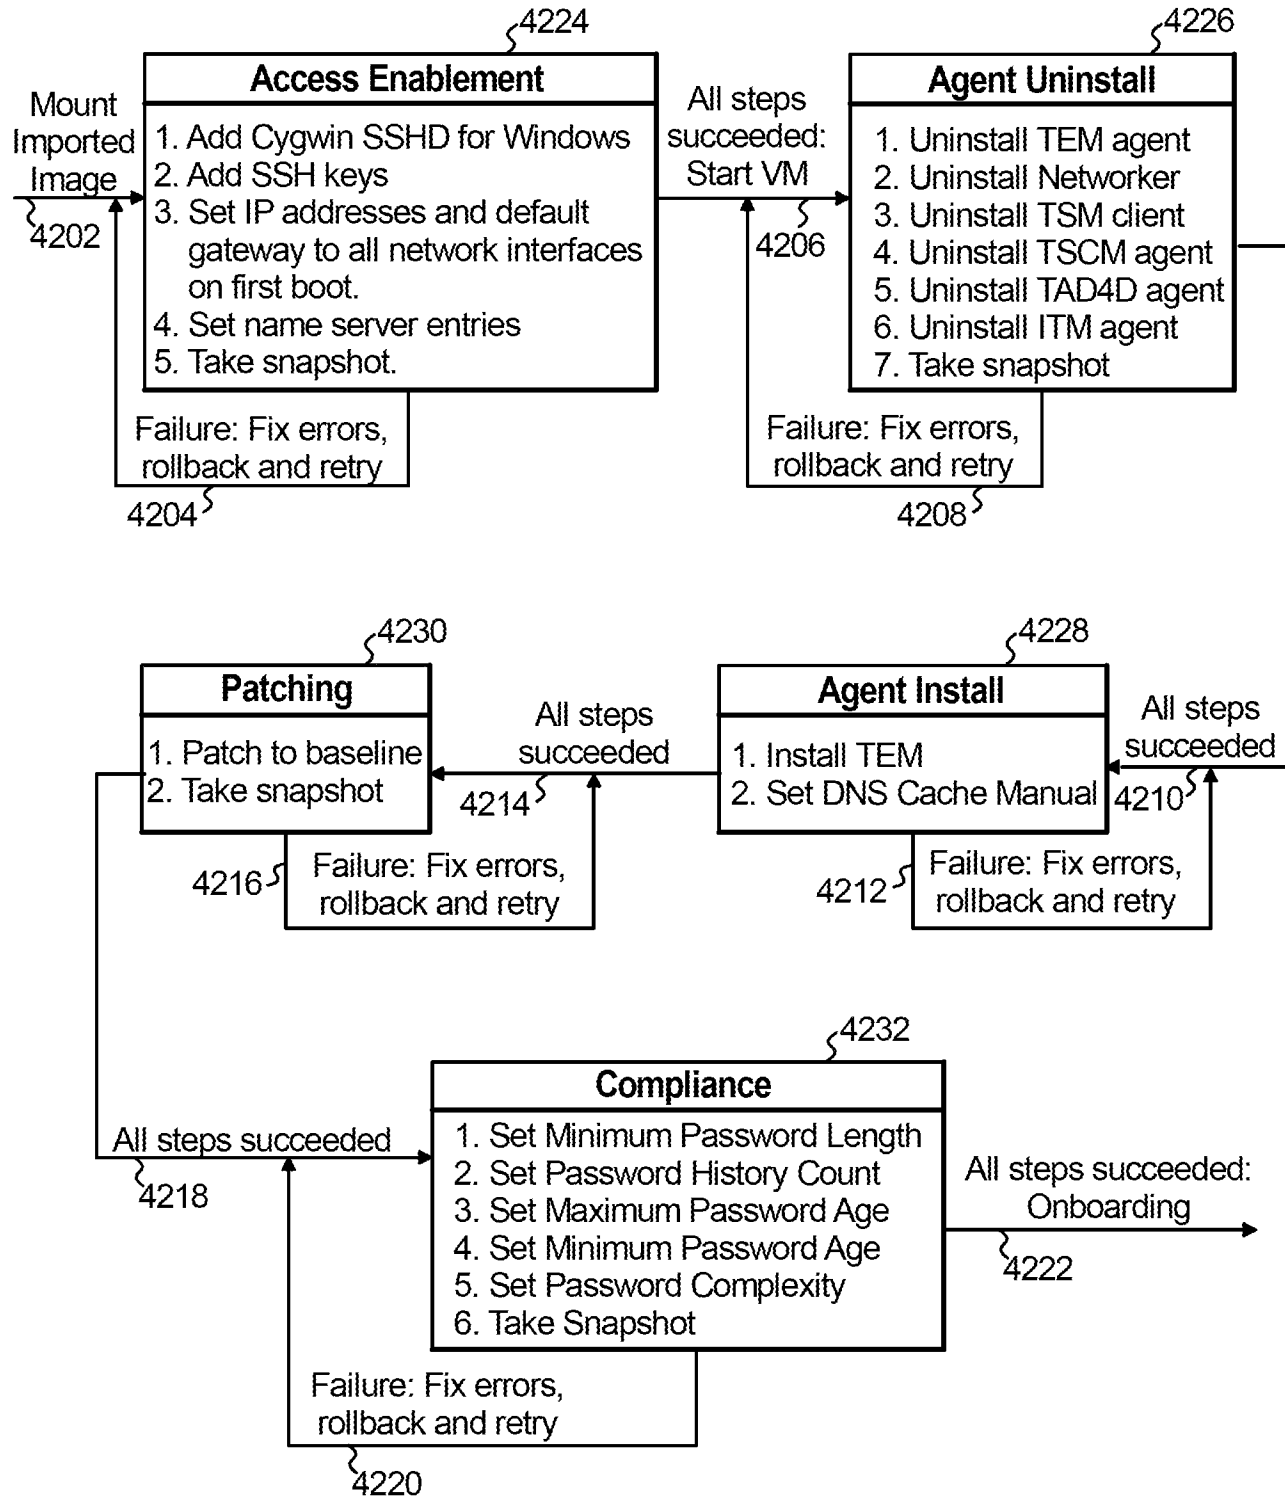


Fig. 42

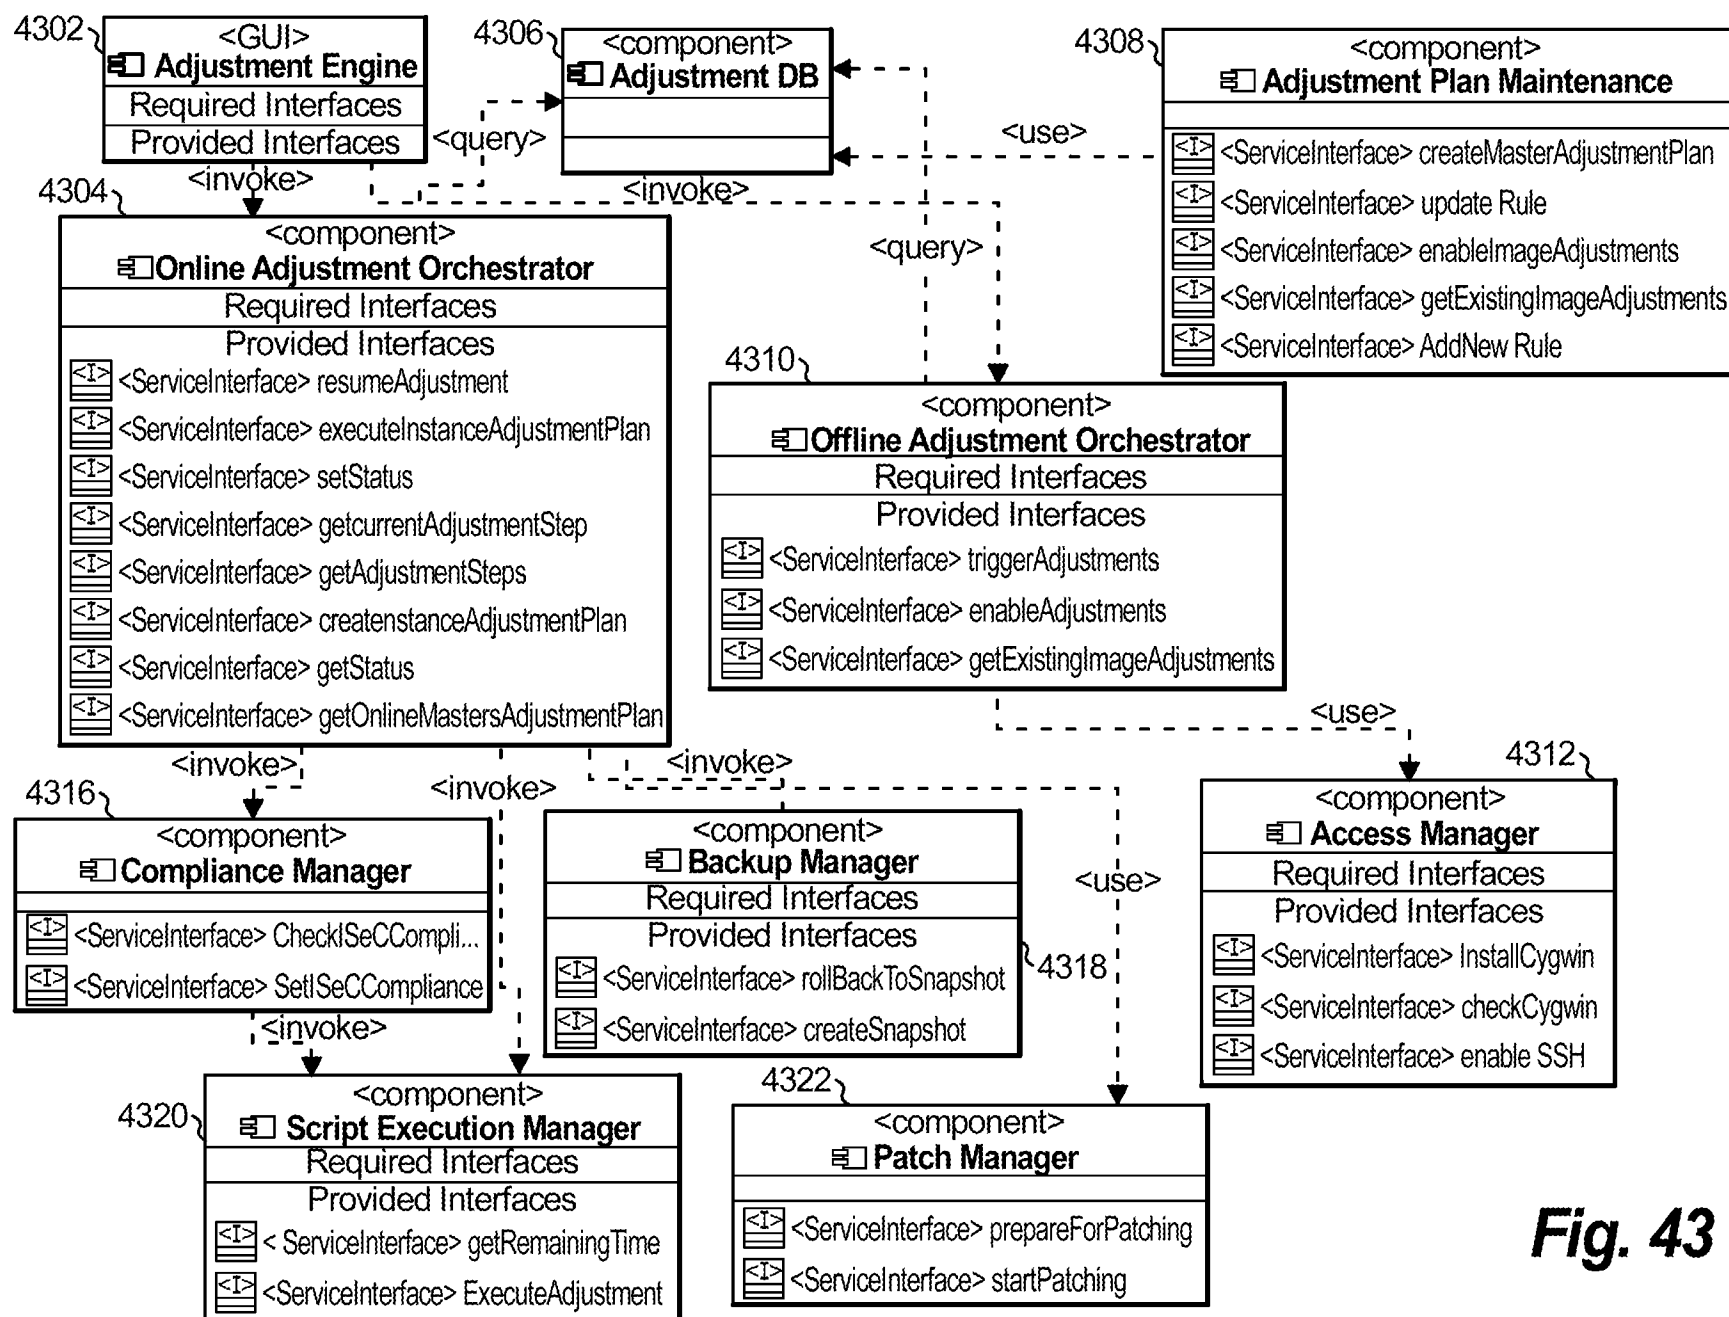
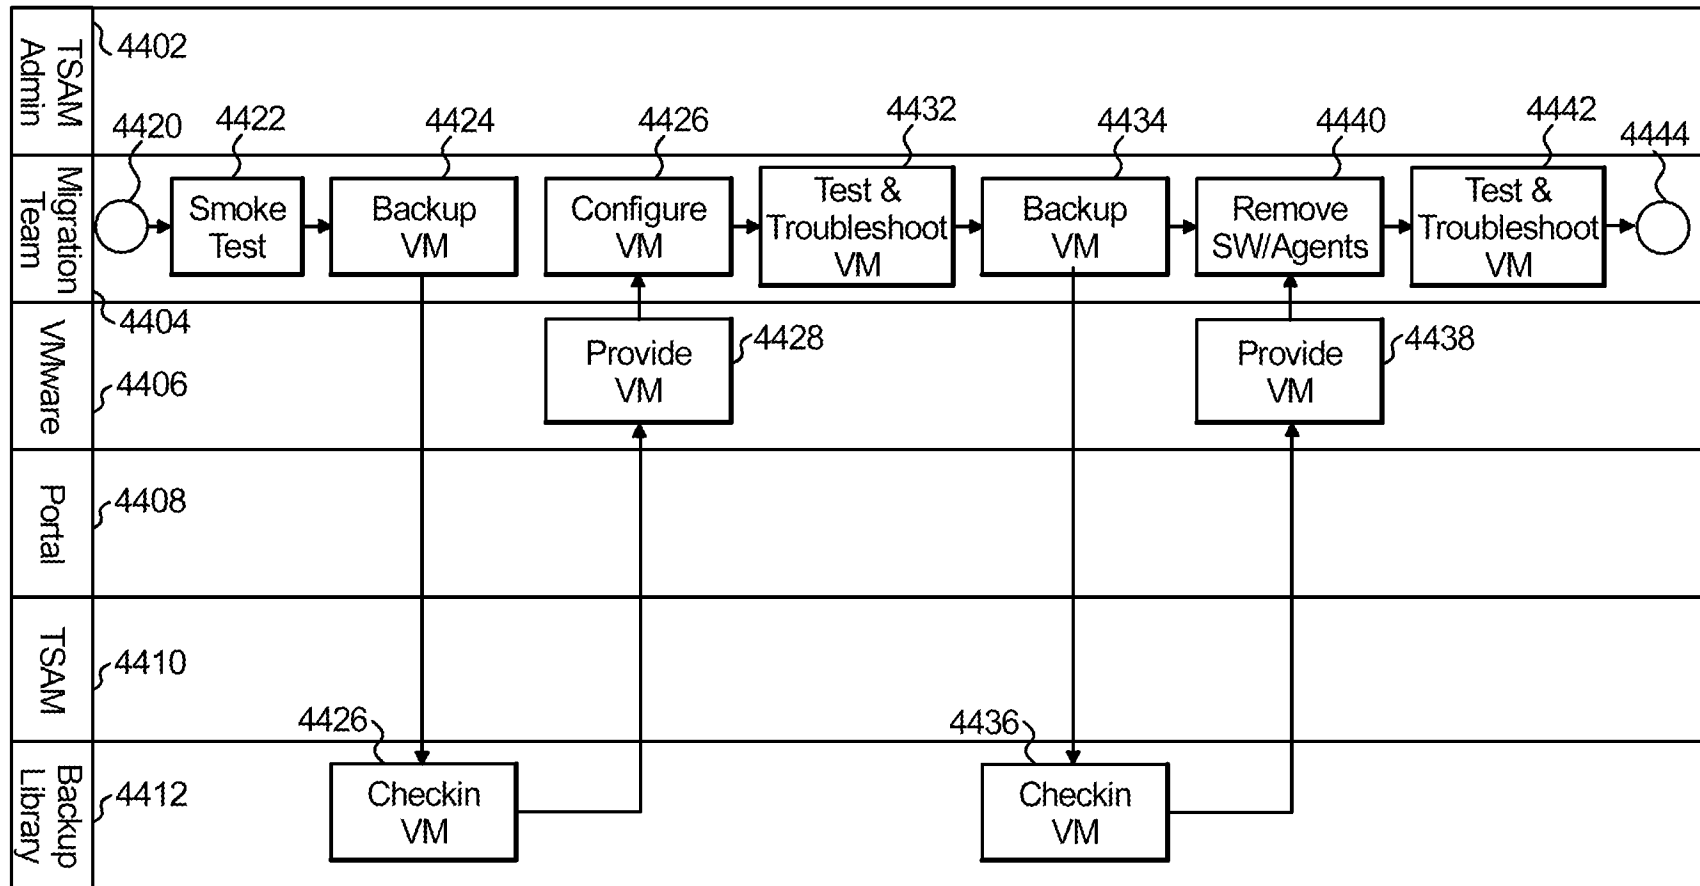


Fig. 43

43/47

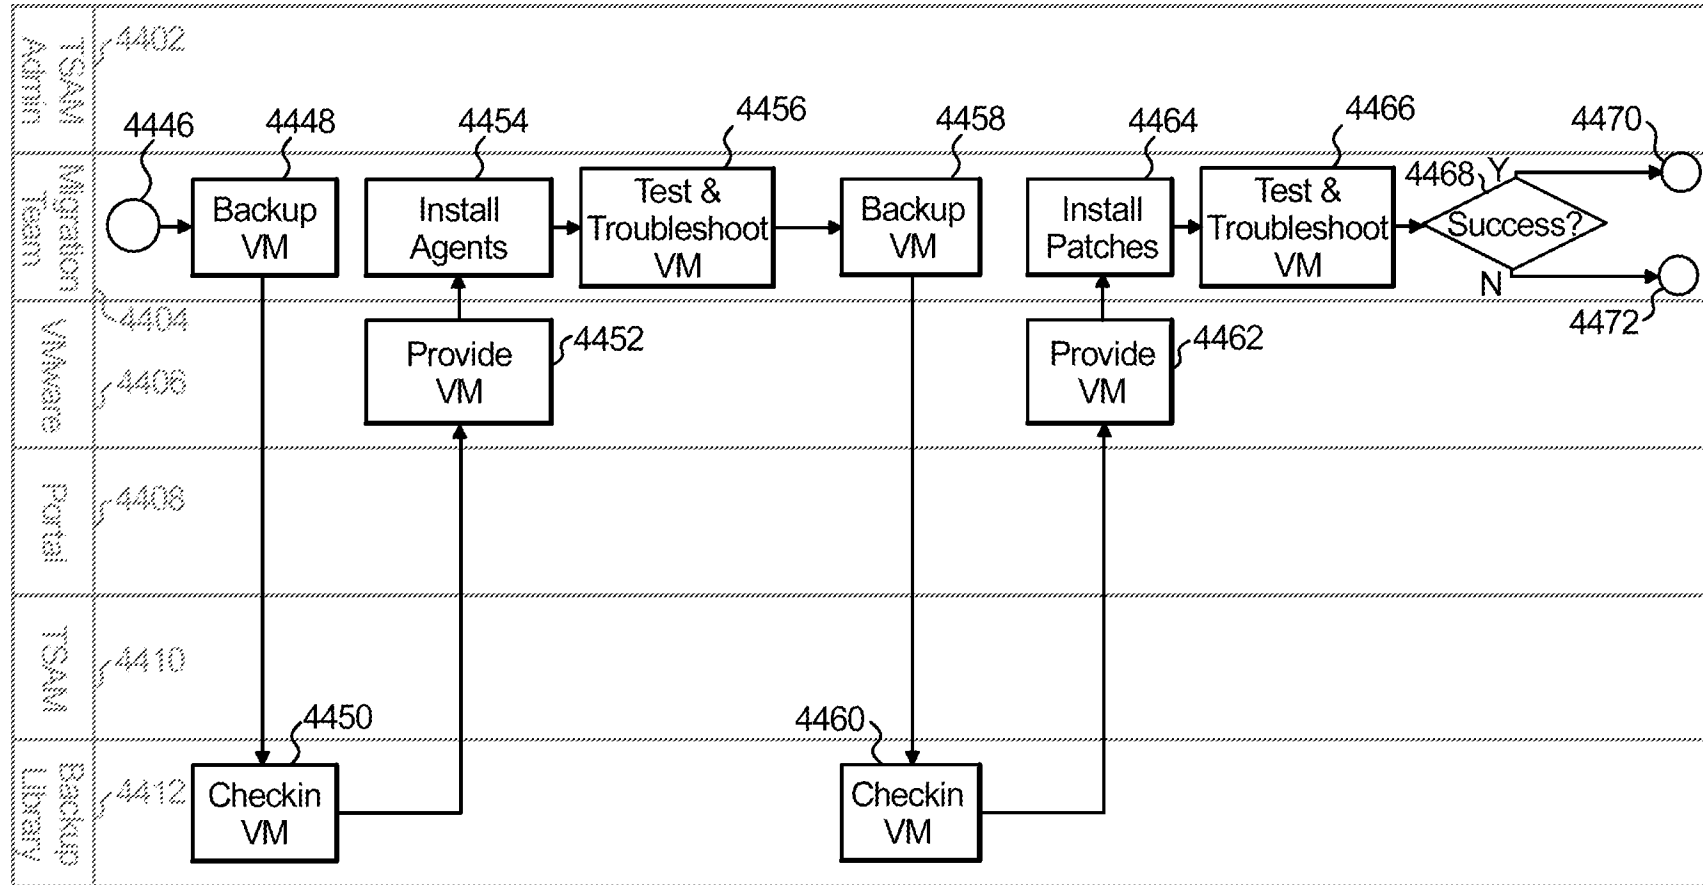
YOR920120450US2



44/47

YOR920120450US2

Fig. 44



45/47

YOR920120450US2

Fig. 45

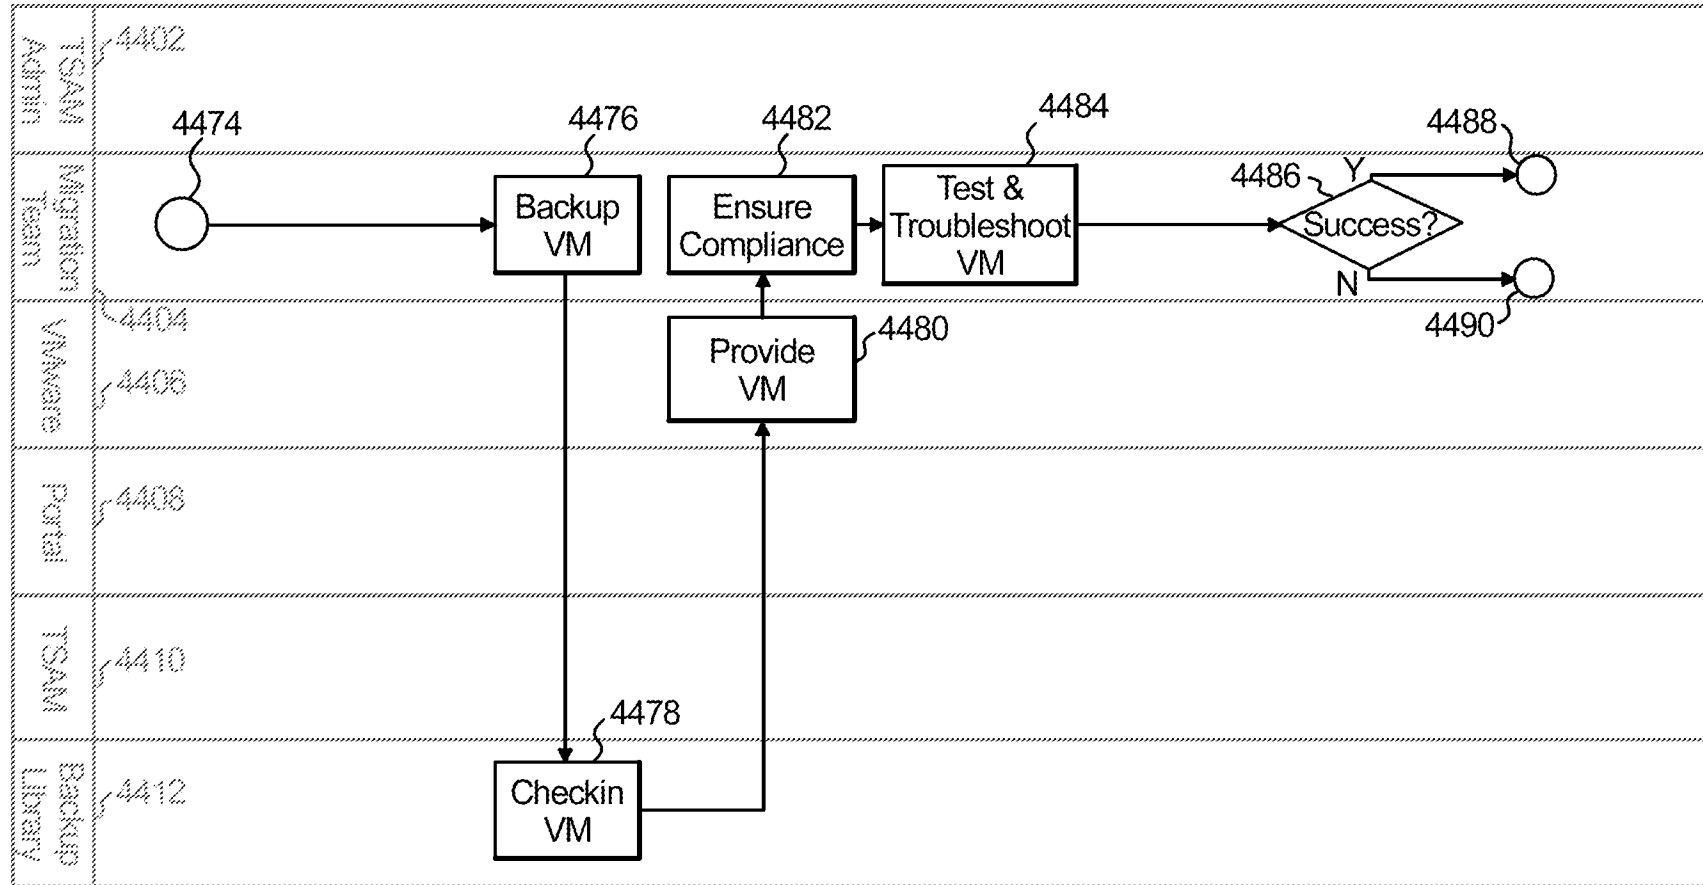


Fig. 46

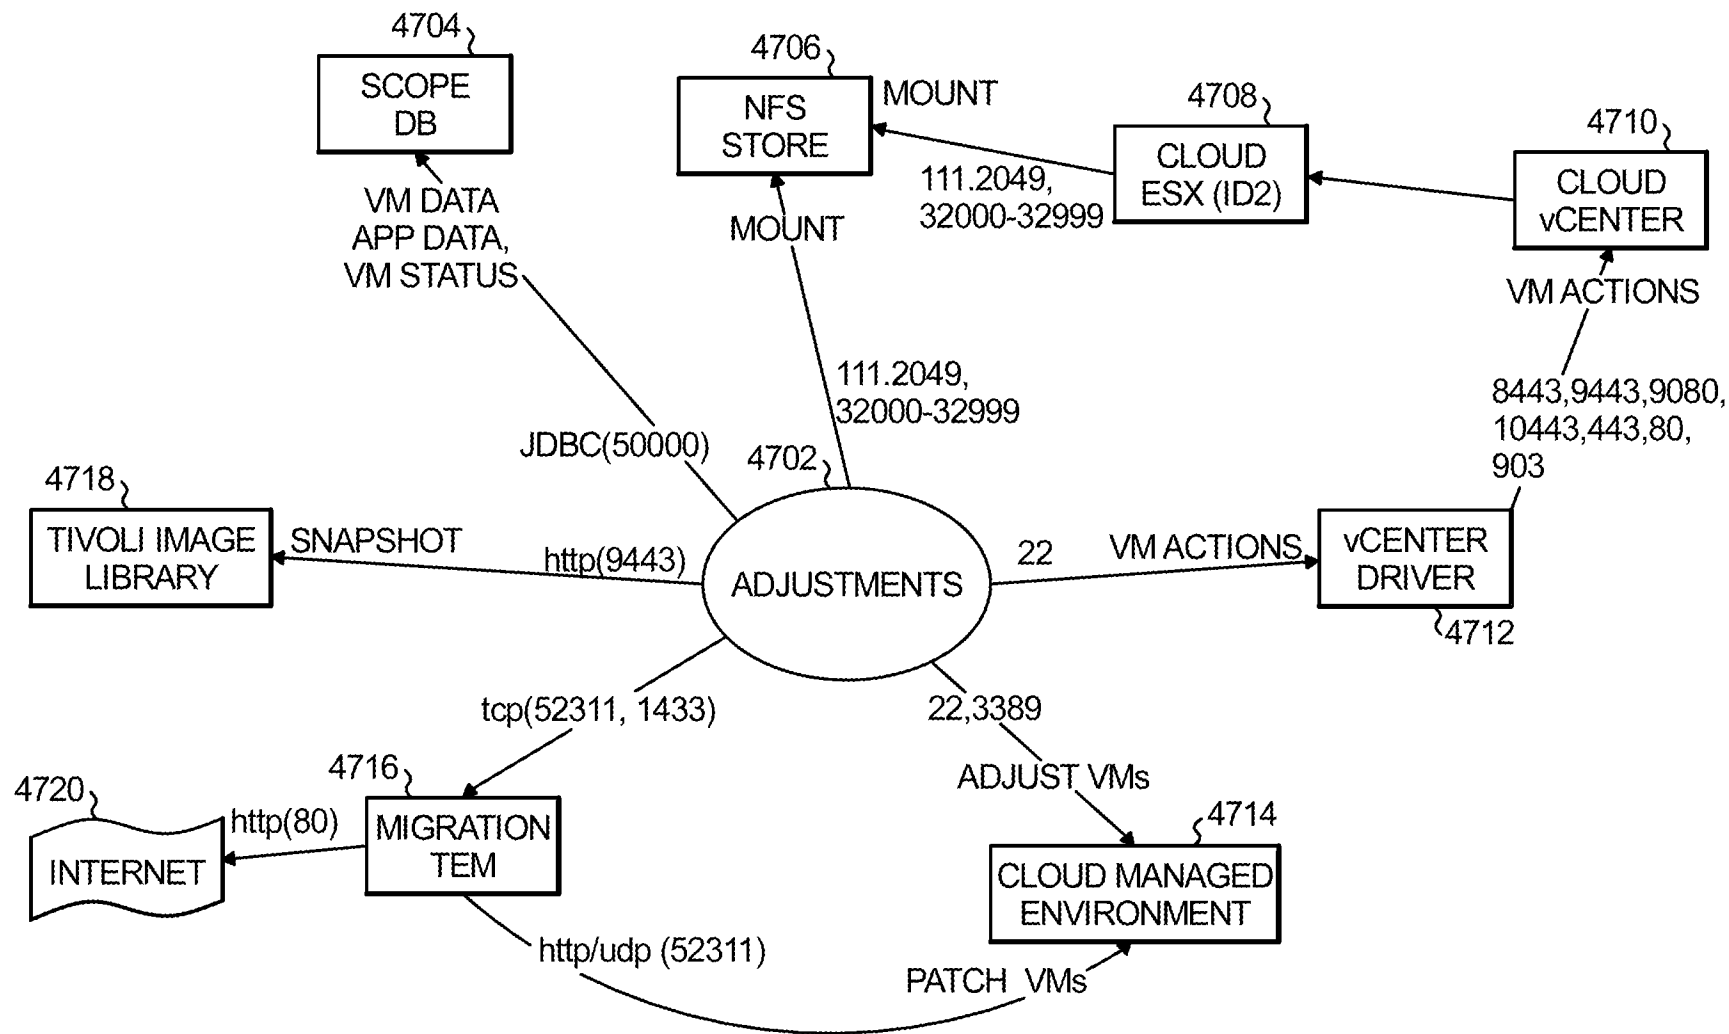


Fig. 47

47/47

YOR920120450US2

MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

Cross-Reference to Related Application

5 This application is a continuation of U.S. patent application no. 13/689,500 filed 11-29-2012, the complete disclosure of which is expressly incorporated herein by reference in its entirety for all purposes.

Field of the Invention

10 The present invention relates to the electrical, electronic and computer arts, and, more particularly, to cloud computing and the like.

Background of the Invention

15 In an Infrastructure as a Service (IaaS) cloud, the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host
20 firewalls).

 Migration is the process of moving from a source environment to a target environment, such as a target IaaS cloud environment. Heretofore, migration analysis has typically only focused on operating systems (OS) and applications, not infrastructure management.

25

Summary of the Invention

 Principles of the invention provide techniques for management infrastructure analysis for cloud migration.

30 In one aspect, an exemplary method includes discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component; and obtaining a description of a target cloud infrastructure having a target management infrastructure. The description includes at least one

mandatory target infrastructure management component. A further step includes analyzing the at least one source infrastructure management component to determine whether at least one conflict exists with the at least one mandatory target infrastructure management component.

5 As used herein, “facilitating” an action includes performing the action, making the action easier, helping to carry the action out, or causing the action to be performed. Thus, by way of example and not limitation, instructions executing on one processor might facilitate an action carried out by instructions executing on a remote processor, by sending appropriate data or commands to cause or aid the action to be performed. For the
10 avoidance of doubt, where an actor facilitates an action by other than performing the action, the action is nevertheless performed by some entity or combination of entities.

One or more embodiments of the invention or elements thereof can be implemented in the form of a computer program product including a computer readable storage medium with computer usable program code for performing the method steps
15 indicated. Furthermore, one or more embodiments of the invention or elements thereof can be implemented in the form of a system (or apparatus) including a memory, and at least one processor that is coupled to the memory and operative to perform exemplary method steps. Yet further, in another aspect, one or more embodiments of the invention or elements thereof can be implemented in the form of means for carrying out one or
20 more of the method steps described herein; the means can include (i) hardware module(s), (ii) software module(s) stored in a computer readable storage medium (or multiple such media) and implemented on a hardware processor, or (iii) a combination of (i) and (ii); any of (i)-(iii) implement the specific techniques set forth herein; the overall apparatus will typically include (ii) or (iii).

25 Techniques of the present invention can provide substantial beneficial technical effects. For example, one or more embodiments provide techniques for management infrastructure analysis for cloud migration with one or more of the following advantages:

- facilitates migration of existing environments to cloud environments having
30 standardized infrastructure management systems and processes (e.g., MIaaS cloud environments),

- improves coverage of existing servers or instances that can be migrated to such cloud environments, i.e., the fraction of servers or instances that can be migrated,
 - decreases the risk of failure of the migration or problems after the migration, and/or
- 5 • where migration is not possible or needs specific preceding measures, decreases the time needed and the risk during the critical actual migration period (where applications are typically shut down) by advance analysis.

10 These and other features and advantages of the present invention will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

Brief Description of the Drawings

15 FIG. 1 depicts a cloud computing node according to an embodiment of the present invention;

FIG. 2 depicts a cloud computing environment according to an embodiment of the present invention;

FIG. 3 depicts abstraction model layers according to an embodiment of the present invention;

20 FIG. 4 presents a high-level flow chart, according to an aspect of the invention;

FIG. 5 presents a detailed flow chart and block diagram, according to an aspect of the invention;

FIG. 6 presents an exemplary system block diagram, according to an aspect of the invention;

25 FIG. 7 presents an exemplary system diagram including system states at significant phases, according to an aspect of the invention;

FIG. 8 shows a current server in a source environment and after migration to a cloud, according to an aspect of the invention;

FIGS. 9 and 10 show alternative migration approaches for the server of FIG. 8;

30 FIG. 11 shows an exemplary provisioning flow, according to an aspect of the invention;

FIG. 12 shows instance capture, in accordance with an embodiment of the invention;

FIG. 13 shows an adoption and adjustment process, in accordance with an embodiment of the invention;

5 FIG. 14 shows an exemplary screen view of a “create server” pop-up, in accordance with an embodiment of the invention;

FIG. 15 shows an exemplary overview of adjustments in a provisioning flow, according to an aspect of the invention;

10 FIG. 16 presents a combined flow chart and block diagram, according to an aspect of the invention;

FIG. 17 depicts a source environment to be migrated to a target environment in accordance with one or more aspects of the invention;

FIG. 18 depicts a target environment to which the source environment of FIG. 5 can be migrated in accordance with one or more aspects of the invention;

15 FIG. 19 shows exemplary phases in management infrastructure analysis for cloud migration in accordance with one or more aspects of the invention;

FIG. 20 shows an exemplary user interface in accordance with one or more aspects of the invention; and

20 FIG. 21 shows non-limiting exemplary results for a small subset of servers in accordance with one or more aspects of the invention;

FIG. 22 presents an exemplary system block diagram, according to an aspect of the invention;

FIG. 23 presents an exemplary snapshot manager, according to an aspect of the invention;

25 FIG. 24 presents an exemplary implementation based on a virtual image library, according to an aspect of the invention;

FIG. 25 presents an exemplary rollback to a long term snapshot, according to an aspect of the invention;

30 FIG. 26 is a flowchart showing steps to compute exact changes made to an image during migration, according to an aspect of the invention;

FIG. 27 is an exemplary process flow, according to an aspect of the invention; and

FIG. 28 is an exemplary software architecture diagram of a snapshot management system, in accordance with an aspect of the invention;

FIG. 29 presents an exemplary system diagram, according to an aspect of the invention;

5 FIG. 30 presents an exemplary method, according to an aspect of the invention;

FIG. 31 presents an exemplary “swap” flow chart, according to an aspect of the invention;

FIG. 32 presents an exemplary “merge virtual resource descriptors” flow chart, according to an aspect of the invention;

10 FIG. 33 presents another exemplary system diagram, according to an aspect of the invention;

FIG. 34 depicts pertinent aspects of virtual machine images and instances;

FIG. 35 shows a first embodiment of a virtual machine resource descriptor, according to an aspect of the invention;

15 FIG. 36 shows a second embodiment of a virtual machine resource descriptor, according to an aspect of the invention;

FIG. 37 is a flow chart of exemplary method steps for updating a target virtual machine descriptor, according to an aspect of the invention;

20 FIG. 38 is a flow chart of exemplary preparatory method steps, according to an aspect of the invention;

FIG. 39 is an exemplary software architecture diagram, according to an aspect of the invention;

FIG. 40 presents a standardization framework, according to an aspect of the invention;

25 FIG. 41 presents a flow of a sample offline adjustment, according to an aspect of the invention;

FIG. 42 presents an exemplary flow diagram, according to an aspect of the invention;

30 FIG. 43 presents an exemplary standardization architecture, according to an aspect of the invention;

FIGS. 44-46 shown an exemplary adjustment phase, according to an aspect of the invention; and

FIG. 47 shows an exemplary architecture, according to an aspect of the invention.

5

Detailed Description of Preferred Embodiments

Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

Characteristics are as follows:

15 On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

5 Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g.,
10 cloud bursting for load balancing between clouds).

A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes.

Referring now to FIG. 1, a schematic of an example of a cloud computing node is
15 shown. Cloud computing node 10 is only one example of a suitable cloud computing node and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, cloud computing node 10 is capable of being implemented and/or performing any of the functionality set forth herein.

20 In cloud computing node 10 there is a computer system/server 12, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server 12 include, but are not limited to, personal computer systems, server
25 computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

30 Computer system/server 12 may be described in the general context of computer system executable instructions, such as program modules, being executed by a computer

system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 12 may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

As shown in FIG. 1, computer system/server 12 in cloud computing node 10 is shown in the form of a general-purpose computing device. The components of computer system/server 12 may include, but are not limited to, one or more processors or processing units 16, a system memory 28, and a bus 18 that couples various system components including system memory 28 to processor 16.

Bus 18 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus.

Computer system/server 12 typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server 12, and it includes both volatile and non-volatile media, removable and non-removable media.

System memory 28 can include computer system readable media in the form of volatile memory, such as random access memory (RAM) 30 and/or cache memory 32. Computer system/server 12 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system 34 can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a

removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to bus 18 by one or more data media interfaces. As will be further depicted and described below, memory 28 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

Program/utility 40, having a set (at least one) of program modules 42, may be stored in memory 28 by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Program modules 42 generally carry out the functions and/or methodologies of embodiments of the invention as described herein.

Computer system/server 12 may also communicate with one or more external devices 14 such as a keyboard, a pointing device, a display 24, etc.; one or more devices that enable a user to interact with computer system/server 12; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server 12 to communicate with one or more other computing devices. Such communication can occur via Input/Output (I/O) interfaces 22. Still yet, computer system/server 12 can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter 20. As depicted, network adapter 20 communicates with the other components of computer system/server 12 via bus 18. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server 12. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

Referring now to FIG. 2, illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 comprises one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate.

Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 2 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

Referring now to FIG. 3, a set of functional abstraction layers provided by cloud computing environment 50 (FIG. 2) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 3 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

Hardware and software layer 60 includes hardware and software components. Examples of hardware components include mainframes, in one example IBM® zSeries® systems; RISC (Reduced Instruction Set Computer) architecture based servers, in one example IBM pSeries® systems; IBM xSeries® systems; IBM BladeCenter® systems; storage devices; networks and networking components. Examples of software components include network application server software, in one example IBM WebSphere® application server software; and database software, in one example IBM DB2® database software. (IBM, zSeries, pSeries, xSeries, BladeCenter, WebSphere, and DB2 are trademarks of International Business Machines Corporation registered in many jurisdictions worldwide).

Virtualization layer 62 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers; virtual storage; virtual networks, including virtual private networks; virtual applications and operating systems; and virtual clients.

In one example, management layer 64 may provide the functions described below. Resource provisioning provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing

environment. Metering and Pricing provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal provides access to the cloud computing environment for consumers and system administrators. Service level management provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

Workloads layer 66 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation; software development and lifecycle management; virtual classroom education delivery; data analytics processing; transaction processing; and mobile desktop.

MIGRATION TO MANAGED CLOUDS

A Hardware-Infrastructure-as-a-Service (HIaaS) cloud provides bare-bones virtual machines as a service. It may also provide an operating system (OS) and even software, but no support is typically provided for the OS or software. A Managed-Infrastructure-as-a-Service (MIaaS) cloud provides full-service virtual machines. The service may, e.g., include OS patching and support for security and compliance of the OS. One significant aspect of MIaaS is simpler management via: standardization to a certain set of catalog images from which instances are generated, automatic linkage of these instances to management tools during deployment, and/or not giving customers administration privileges on the OS level, so that the operating systems on the instances remain as the cloud administrators configure them.

An MIaaS cloud may not naturally come with import or registration features for foreign instances, inasmuch as a significant aspect of MIaaS is simpler management via: standardization to a certain set of catalog images from which instances are generated; automatic linkage of these instances to management tools during deployment; and,

typically, not giving customers administration privileges on the operating system (OS) level, so that the operating systems on the instances remain as the cloud administrators configure them. Thus, in migrating to an MIaaS cloud, it is typically not feasible to simply use P2V (physical-to-virtual) conversion on source instances, or to directly copy
5 already virtualized instances, and then expect them to be run on cloud hypervisors. This is so because they would not fulfill the aforementioned criteria through which simpler management is achieved in the MIaaS cloud, and thus they would not be acceptable to the management of the MIaaS cloud.

Furthermore, with regard to MIaaS clouds, the standard registration (i.e., making
10 a new instance known to the general IaaS cloud management systems as well as the specific management systems of an MIaaS cloud) is typically inbuilt into the provisioning from catalog images. Advantageously, one or more embodiments provide a new registration procedure wherein foreign instances can be accommodated within migration to an MIaaS cloud. Indeed, one or more embodiments advantageously provide a system
15 and method for rapid migration into MIaaS (and more generally, IaaS) clouds. The method includes transfer of the instance to the cloud in image form; adjustments of the instance (running or in image form) to cloud standards; and registration of the instance into cloud OSS and BSS systems (Operation and Business Support Systems). Optional additional steps address prior analysis, testing and handling failures, and/or starting and
20 ending change windows and actual downtimes in order to minimize risk and outages.

One or more embodiments advantageously provide a systematic (and even automated) method of rapidly migrating customer instances into MIaaS clouds that does not involve a reinstall procedure. One or more embodiments are useful with MIaaS cloud migration and are capable of physical-to-virtual style instance importation.

25 One or more embodiments enhance the ability to standardize both infrastructure *per se* and the management of such infrastructure. It is believed that such standardization will in turn allow reduction in IT operational costs and permit further automation. One or more embodiments provide techniques for migration into an MIaaS cloud which are significantly less expensive than techniques that require reinstall.

30 As noted above, one or more embodiments provide a migration technique with one or more of the following advantages:

- significant coverage (percentage of instances and then workloads where the method can be used),
- low cost, which in particular is bounded by manual work needed (in particular compared with reinstall methods)
- short migration times (as change windows / permitted outages are typically short),
- low risk (e.g., of application outages beyond the planned ones), and/or
- predictability (i.e., workloads chosen for this migration technique will likely succeed in it).

10

With attention now to FIG. 4, begin with an instance 401 in an existing customer environment 402 (seen in FIG. 5) which it is desired to migrate. In step 404, transfer the instance to the cloud in image form. If this is not successful, retain the original customer version 401 and re-plan the migration; for example, retain it in its original environment, or perform a classical migration into a service provider environment where it can remain a physical instance, or make significant changes on the instance so that it becomes virtualizable. On the other hand, if step 404 is successful, in step 406, carry out adjustments of the instance (in running and/or in image form) to ensure that it complies with cloud standards. If this is not successful, retain the original customer version 401 and re-plan the migration; for example, virtualize it in the customer environment or a service provider environment without cloud standards, or an HIaaS cloud, or perform a more complex reinstall migration wherein individual software components are newly installed and customer data are transferred individually rather than within the entire image. On the other hand, if step 406 is successful, in step 408, carry out registration of the instance into the cloud OSS and BSS systems. If step 408 is not successful, retain the original customer version 401 and re-plan the migration; for example, by the same approaches as when steps 404 or 406 failed. On the other hand, if step 408 is successful, the result is successful migration to the MIaaS cloud 410.

Note that initial sub-steps for the cloud registration may take place before the adjustments or may be interleaved with the adjustments. However, from a logical

30

standpoint, the final registration, i.e., final acceptance of the imported instance as a cloud-managed instance, should be after the end of the adjustments.

Turning now to FIG. 5, a non-limiting exemplary method is depicted therein with greater detail. A customer environment 402 is connected to an MIaaS cloud environment 410 via a wide area network (WAN) 514 or the like. Initially, a discovery process 518 is carried out to determine both the physical 520 and virtual 522 instances and their configurations in the customer environment 402. Analysis and planning is then carried out in step 524. If the results are unfavorable, other methods, such as physical-to-physical (P2P) migration, application reinstallation, retention of the legacy systems, and the like are pursued, as shown at 526. Also, if either of steps 518, 524 indicate that remediation (i.e., small changes to the instances to make them compatible with the cloud) is required, the same is carried out at 516 and then process flow returns to step 518. On the other hand, if step 524 indicates that migration using one or more techniques disclosed herein is feasible, flow proceeds to baseline test and backup, in step 528. In step 528, as shown at 530, a set of test cases are run to prove that the source system fulfills criteria that are desired to be maintained in the target environment, e.g., that it still tests correct under all of its functional test cases or performance claims. This is done to ensure that any errors already present are fixed before the migration process. Furthermore, a back-up process is carried out, as shown at 532, to permit recovery in case of any issues encountered in the migration process. Preferably, before finalizing the backup 532, applications on the physical or virtual instance 520 or 522 are stopped so that no further changes take place in the customer environment under conditions such that they would not be copied to the MIaaS cloud environment. Processing then proceeds to step 534, wherein the instances to be migrated are captured. One or more illustrative embodiments focus on specific aspects of migrating to an MIaaS cloud, which are largely per-instance. As to the overall migration, it is typically done in waves of multiple instances each, e.g., per weekend, attempting to keep workloads or interacting instances within one wave. Such aspects have been disclosed, e.g., in Athey et al., System and method for object migration using waves, United States Patent Application Publication 20110213883, September 1, 2011; and Devarakonda et al., Total cost-based migration waves planning, United States Patent Application Publication 20120109844, May 3, 2012.

YOR920120450US2

2012, both of which are expressly incorporated herein by reference in their entireties for all purposes.

As seen at 536, this instance capture step can include, for example, both physical-to-virtual (P2V) and virtual-to-virtual (V2V) techniques with one or more suitable tools.

5 One non-limiting example of a suitable tool is PlateSpin® Migrate, a physical/virtual conversion tool for fast and efficient P2V (more broadly, anywhere-to-anywhere) migrations; it is available from NetIQ Corporation, Houston, Texas, USA; another is VMware vCenter Converter, available from VMware, Inc., Palo Alto, CA, USA. In one or more non-limiting exemplary embodiments, the end result of step 534 is a Virtual
10 Machine Disk Format (VMDK) file.

It is to be emphasized that many product names are mentioned herein; these are intended as examples for the skilled artisan and to convey the applicants' understanding of the best mode. They are not intended to limit the claims unless expressly recited therein, but rather are to be considered exemplary of corresponding generic software
15 products; e.g., PlateSpin® Migrate is broadly representative of physical/virtual conversion tools.

As seen in step 538, the captured instances are then transported into the cloud location 410 via network 514. Data 544 outside the boot disk may be transported separately from the aforementioned vmdk file (seen at 542), in particular if it is large and
20 the data transport might start earlier. As noted at 540, the instances and data are transported through network 514 using appropriate tools to control the transfer. Non-limiting examples of such tools include the aforementioned PLATESPIN tool, as well as the Softek Transparent Data Migration Facility (TDMF®) tool for data (registered mark of International Business Machines Corporation, Armonk, NY, USA). Data 544 is
25 typically not affected by special aspects of MIaaS cloud, i.e., it can be migrated and linked back with the vmdk in usual ways. Thus further details about the data 544 are omitted from the figure to avoid clutter.

In step 546, functional testing is carried out on the transported instance(s) 554 that was restarted from the image on a cloud hypervisor. As seen at 547, this can include, for
30 example, playing a set or subset of test cases; this can be repeated several times between adjustments. Also, it should be noted that upon initial arrival in cloud environment 410,

the transported instances reside within MIaaS cloud landing zone 572. If functional testing is successful at this time, processing flow proceeds to instance adjustment and adoption 556. This is a significant step to modify the instance so that it can run in the standard MIaaS environment. After all adjustments, and possibly between certain adjustment steps, functional testing is repeated, as shown by the backwards error from Step 556 to 546. After each successful functional test, instances can be backed up (i.e., a snapshot 552 of the instance can be taken) as shown at step 550 (e.g., as vmdk files in an efficient instance repository 552) for later reference, as a most recent correct-seeming state of the instance. If functional testing is successful after all adjustments 556, the transported instance(s) are instantiated on cloud hypervisor 558 within MIaaS cloud production zone 574. On the other hand, if functional testing is not successful in one of its repetitions, processing flow proceeds to remediation step 548. The remediation process can proceed via a number of avenues. Typically, error messages will be analyzed and related to the latest changes. These latest changes are the virtualization and new hypervisor in the first execution of the functional testing 546. In subsequent executions, they are the changes by the adjustments since the previous execution of the functional. For this, it may be useful to refer to a snapshot 552 of the instance before these last changes. If remediation of these errors is not successful (which might be determined by a rerun of the functional test), flow returns to step 528, i.e., the migration of this instance is stopped (at least for the moment) and the backup is used to restore the source instance 520 or 522 in the customer environment 402. When remediation is successful, the processing proceeds as if that functional test had succeeded at once, i.e., with the first or next adjustment or, if all adjustments are done, with instantiating the instance on cloud hypervisor 558 within MIaaS cloud production zone 574.

Within the instance adjustment and adoption step 556 the instances on hypervisor 554 are adjusted to cloud delivery standards and adopted into the cloud BSS and OSS, as seen at 568. As seen at 570, this process employs provisioning flows extended by adjustment subflows. For example, the standard provisioning flow of the MIaaS (i.e., for instances chosen from the cloud catalog, rather than migrated) may register a new instance in an asset management system, a monitoring system, and start billing and accounting for it. This functionality can be reused from the standard provisioning flow.

On the other hand, the standard provisioning flow may not install a certain monitoring agent that the MIaaS cloud requires on the normal instances because they would be preinstalled in the cloud catalog images. Hence installing this agent would be part of the special adjustment flows. Similarly, an update of the security patches to the level that the cloud catalog images have may be an extra step in the adjustment flows. When the functional testing after all adjustment steps finally has successfully completed, as noted, instantiation occurs on cloud hypervisor 558 within production zone 574.

In one or more embodiments, the landing zone is a special zone added for migration to an MIaaS cloud, because when the instances are first started on a hypervisor the MIaaS cloud environment (see item 554) in order to test and adjust, they do not yet fulfill cloud standards and thus the normal cloud management cannot handle them. They may also not fulfill the security standards that the OS management in the cloud provides and assumes. For this reason, one may physically separate the landing zone 572 and the production zone 574, at least by using different servers to host the hypervisors, and separate them by firewalls that only let the controlled information through as indicated by the two arrows between the zones. Also the storage systems may be separate. Another advantage of physical separation is that the cloud management systems such as capacity and performance management then don't have to deal with hypervisor(s) that are partially filled with normal, managed cloud instances and partially with as yet unmanageable imported instances. However, it is also conceivable to use logical separation, i.e., trust the cloud hypervisors to not let the potentially insecure imported images impact the others, and to extend the cloud management systems to deal with partitioned hypervisors.

Once the pertinent migrated instances are instantiated on cloud hypervisor 558 in production zone 574, user acceptance testing 560 may be conducted, for example, to verify that performance is equivalent to the baseline. If not so, remediation is carried out at 562, to correct the instances on hypervisor 558. This remediation process can make use of instance repository 552 or alternate techniques 526 as appropriate; if this happens, the instance should be removed from the production zone 574 again until it is remediated. The test cases in the user acceptance test 560 should preferably be identical to those in the baseline test 528. They often cover multiple instances together if the user-relevant function is performed by multiple instances, e.g., a web server, an application server, and

a database. In contrast, the functional test may be per instance only. When testing 560 is successful, cut-over is carried out in step 564, taking into account any customer domain name server (DNS) changes or the like, as seen at 565. Upon completion of cut-over, as seen at 566, the migrated instances operate in the MIaaS cloud production zone in a business as usual (BAU) manner.

In FIG. 5, it will be appreciated that in one or more exemplary embodiments, steps 516, 528, and 560 are in the scope of the customer while steps 518, 524, 526, 534, 538, 550, 556, and 564 are in the scope of the cloud services provider. Steps 546, 548, and 562 are of a mixed nature between both parties.

FIG. 6 shows an exemplary system diagram. Customer environment 402 includes one or more source instances 684 (corresponding to the initial physical or virtual instances 520 and 522 in FIG. 5) and a transfer agent 682. Network 514 provides connectivity between customer environment 402 and MIaaS cloud environment 410. Environment 410 includes landing zone 572, production zone 574, and management zone 680. Transferred instances 686 reside within landing zone 572 (element 554 represents an instance running on a hypervisor). Final target instances 688 reside within production zone 574 (element 558 represents an instance running on a hypervisor). Cloud management zone 680 includes a number of management components. Transfer core component 690 carries out instance transfer of source instance 684 to transferred instance 686, working in conjunction with transfer agent 682. Adjustment component 692 orchestrates the adjustment process discussed with regard to landing zone 572 in FIG. 5. Registration component 694 registers the adjusted instance into the normal management system of the MIaaS cloud, here shown by the BSS and OSS 696, 698. Cloud OSS and BSS systems (Operation and Business Support Systems) 698, 696 provide input for step 556, as seen at 568. The normal function of OSS and BSS in an MIaaS cloud is to manage instances in the production zone 574. However, the landing zone 572 also needs at least a minimal OSS, e.g., for capacity management and thus finding a suitable place (server and storage with enough capacity) when a source instance 684 is coming in. If the landing zone is physically separated from the production zone, then at a more detailed level the OSS may be separated too. In any case, the adjustment component 692 and registration component 694 prepare the instance for reception into the OSS 698 so there

is at least an abstract connection in the sense of them needing knowledge about the OSS requirements. A BSS for the landing zone is typically not necessary because in the landing zone, the instances are controlled by the migration process, and not yet visible to the customer or end user.

5 FIG. 7 shows system states at significant points in time (i.e., during significant phases); namely, initial state 707, after transport 709, and after cutover (cloud BAU) 711. In state 707, the customer's existing management tools 703 are managing both the physical 520 and virtual 522 instances in the customer environment 402. The same are stored in storage 705 (e.g. as described with respect to backup step 528 and 532, but also
10 the plain external storage of the instances if that is outside the servers, e.g., on a storage-area network or a networked file system). The instances are captured. As discussed above with regard to 536, this instance capture can include, for example, both physical-to-virtual (P2V) and virtual-to-virtual (V2V) techniques with one or more suitable tools; transport can be via WAN 514 (which could include, for example, the final virtual private
15 network (VPN) (used for communication between the customer and cloud environment during normal operation after migration is complete), or a temporary line) or by transfer of physical media. In one embodiment, cloud environment 410 includes cloud provider hardware and hypervisor 554 in cloud landing zone 572, cloud provider hardware and hypervisor 558 in cloud production zone 574, and shared storage 701. This embodiment
20 is a compromise between full physical separation of landing zone and production zone as described above, where even the storage would be separate, and solely logical separation, where even the hypervisors would be shared and the zone separation would be enforced by the hypervisors – both of those are also among the possible embodiments. One benefit of this mixed embodiment is that the final transfer of the instance from landing zone to
25 production zone does not need to copy the data; only the ownership of the storage volume(s) belonging to this instance is changed to the production zone.

 After transport, as seen at 709, physical and virtual instances 520, 522 are now both virtualized on cloud provider hardware and hypervisor 554 in the landing zone 572, as seen at 720, 722, respectively, and there the adjustments and testing take place.
30 Migrated data is stored in shared storage 701. Preferably, the source instances 520, 522 have been shut down before the transfer, as explained above with respect to the backup

528. One extreme case is that they are not running at all (physical server switched off for 520; instance removed from hypervisor for 522), but typically one would only stop the services on them (so that no changes are made that would be lost when business-as-usual resumes in the cloud production zone). If read-only services can be distinguished, e.g.,
5 browsing on informational web pages, those can keep running. Furthermore, one would typically retain the instances 520 and 522 still in the customer management tools in case migration fails, so that operation in the customer environment could resume quickly.

After cut-over, during the cloud BAU (business as usual) state, as seen at 711, physical, virtual, and storage resources 520, 522, in the customer environment 711 are
10 shut down, and storage 705 used by them is released. Meanwhile, any required adjustments and remediation have been completed, as described with regard to FIG. 5, physical and virtual instances 520, 522 are now both virtualized on cloud provider hardware and hypervisor 558 in the cloud production zone, as seen at 724, 726, respectively. Migrated data is still stored in shared storage 701 in this embodiment of
15 partial physical separation of landing zone and production zone. With full physical separation, there would be two different storage systems, one in the landing zone and one in the production zone, and the data would reside on the former in phase 709 and on the latter in phase 711. BAU cloud processing occurs under control of cloud OSS and BSS systems (Operation and Business Support Systems) 698, 696 seen at 568. The customer
20 management tools 703 may be provided with an interface to production zone 574. However, this typically only concerns those specific tools that manage on the application layer, while tools that perform OS-level functions (such as OS performance management) have been replaced by the cloud OSS in 568. For example, the instance (initially 520 or 522, then 720 or 722, finally 724 or 726 may contain a database, and the customer may
25 have a database management tool as part of 703. An MIaaS cloud (in contrast to a PaaS cloud) typically does not perform database management. So the customer database management tool gets a link to this database on instance 724 or 726. This approach uses the normal way in which customers of the cloud link to their instances of the cloud, e.g., a VPN; it is not in conflict with any zoning or administration constraints of the MIaaS
30 cloud. It should be noted that FIG. 7 shows an exemplary time of one hour to two days to move from initial state 707 to “after transport” state 709 and an exemplary time of two

hours to one day to move from “after transport” state 709 to “after cutover cloud BAU” state 711. These values are intended to provide helpful example for the skilled artisan but are not intended to be limiting.

One or more embodiments thus provide a method for migrating an instance into a cloud, including transferring the instance to the cloud in image form; adjusting of the instance to cloud standards; and registration of the instance into cloud management systems. In some cases, registration is done by a variant of the cloud provisioning method in which picking up a catalog image is replaced by picking up a transferred instance. In some embodiments, cloud management systems include an operational support system and a business support system. In some embodiments, cloud standards include one or more of security standards, infrastructure standards, patch management standards, and infrastructure management tool standards.

Furthermore, in one or more embodiments, adjustment is done by a workflow defined in a workflow engine. In some cases, one or more adjustments are done in the workflow by special tools such as a patch management tool.

In some embodiments, before the transfer step, analysis is carried out on the instance to determine whether it is suitable for the cloud and the given migration method, and only suitable instances are considered further. Furthermore, in some embodiments, tests are carried out at one or more of the times before the transport, after the transport, between different adjustments, after adjustments, and after migration, and remediation and/or back-out occur if one or more of the tests fail. In some cases, multiple instances are treated as a whole in one or more of the decisions to migrate (i.e., whether suitable for cloud as mentioned earlier in this paragraph) and/or in the tests and back-out decisions.

Furthermore, one or more embodiments include discovery, actual handling of failures, starting and ending change windows and actual downtimes, detailed interleaving of adjustment and registration, and/or the cut-over step.

Certain terms employed herein will now be defined:

IaaS clouds: Infrastructure-as-a-Service is a common term for clouds that primarily offer their users virtual machines (VMs), rather than also software on the VMs

(referred to as PaaS, platform-as-a-service) or even software without access to the VMs (referred to as Software-as-a-Service) or business processes (BPaaS).

The following subdivision of IaaS clouds are introduced herein:

5

HIaaS cloud: A Hardware-Infrastructure-as-a-Service cloud provides bare-bone virtual machines as a service. It may also provide an operating system (OS) and even software, but typically no support is provided for the OS or software (e.g., Amazon Elastic Compute Cloud (Amazon EC2) available from Amazon Web Services LLC
10 Seattle, Washington, USA).

MIaaS Clouds: A Managed-Infrastructure-as-a-Service cloud provides full-service virtual machines. The service may, e.g., include OS patching and support for security and compliance of the OS (e.g., IBM SmartCloud Enterprise+ also known as IBM SCE+,
15 cloud environment available from International Business Machines Corporation, Armonk, NY, USA)

Instance: An operating system instance together with all software running on this operating system. It may be physical (i.e., directly running on a server) or virtual (i.e.,
20 already running on a hypervisor).

Source instance: Instance as it is running on the source side, before migration.

Image: File representation of an instance.

25

Catalog image: Image in a cloud catalog, to be used if new instances are created in the cloud from scratch, rather than by rapid migration.

Provisioning: This is the standard mechanism of ordering VM instances in an IaaS
30 - the user orders a VM from the catalog, and the corresponding actual image is instantiated into a running instance.

Reinstall migration: From the point of view of an MIaaS cloud, the standard way of moving an application into same is to first provision a catalog image, and then to reinstall necessary software components, source code, configurations, and data into it.

5 However, this is typically a very time-consuming and costly process.

Image-based migration: Image-based migration uses customer instances for migration of the application to the cloud. In classic virtualization, image-based migration is a standard (P2V) and can be very fast. However, because the image generated directly
10 from the source instance does not enable the management part of the MIaaS cloud, heretofore, image-based migration has not been feasible with MIaaS clouds.

Rapid migration to MIaaS cloud: An extension of image-based migration, in accordance with one or more embodiments, which can handle the challenges of image-
15 based migration, in an MIaaS cloud.

Referring now to FIG. 8, with regard to rapid migration, the skilled artisan will appreciate that there are typically many things on a server 802; for example, middleware or other off-the-shelf software (MW) 804, 806; Infrastructure software (“Infra”) 808 such
20 as monitoring and provisioning agents; custom code 810; and/or scripts 812, e.g., for backup, scheduled cleanup tasks, data transfers, and the like. Of course, the server includes hardware 820 and an operating system (OS) 822 (typically with registry and users (not separately numbered to avoid clutter)). As seen at 824, the current server may be a simple physical server with an OS, data, and software, or may already be virtualized
25 via a suitable hypervisor. Middleware such as 804, 806 typically has configurations 814, data 816 (e.g., databases), code 818 (e.g., SQL scripts), and is often associated with other scripts (such as database administrator scripts). Most of these do not need to be in standard locations. We have found that it is typically quite difficult to automatically find everything that is related to the running instances even of very common standard
30 software such as web servers and databases.

In one or more embodiments, rapid migration in accordance with aspects of the invention is believed to be particularly suitable for an infrastructure-only migration, and for unchanged major OS version, i.e., when the goal is to get to an MIaaS cloud, but not to make other major changes. As seen at 802', when the "same" server is moved into the cloud, it now runs on cloud hardware 830 and hypervisor 832 is employed. Other than that, there are limited changes:

- small updates in the operating system (now OS' 822') – namely, different drivers, cloud provider user IDs, IP address
- somewhat different infrastructure software (cloud management tools – now Infra' 808')

One or more embodiments make use of the insight that it is advantageous to update the few things that change, rather than individually pick up all those that don't change, with a risk of forgetting something. Referring now to FIG. 9, the OS 822; MW 804 (including its configurations 814), 806; infrastructure 808, code 810, scripts 812, code 818, and data 816 are kept together, moved into the cloud, and then small changes are made to OS 822' and infra' 808'. Advantageously, in this approach, if anything does not work in a change, it is most likely only the cloud service provider-side infrastructure management (as that gets newly added), not the application (which remains essentially untouched). Of course, this assumes in the first place that the application can in fact be virtualized. The skilled artisan will of course appreciate, given the description herein, that extra care should be taken, e.g., that the new resources of the server additionally support the cloud management agents.

FIG. 10 shows an alternative approach (reinstall/re-platform) that does not employ the aforementioned rapid migration technique. This is what the MIaaS provider would typically expect its users to do, but is often not feasible in terms of time and money. In this figure, after provisioning an instance with an operating system OS' 808' similar to the original OS 808 from the cloud catalog, and freshly installing the off-the-shelf software, essentially everything is moved one by one. There is some current tooling available along these lines, but it is only available for a small subset of possible off-the-

shelf software, and almost no tooling is available to move extras like scripts, code, etc., nor even just to list it all for subsequent transfer. It is also not in any way an easy option to move “everything else” after installation – OS, applications, and configs / data / custom code are highly linked, both in file systems and in OS data structures such as registries, in particular on WINDOWS.

Consider now core options of interleaving adjustments and registration in rapid migration, i.e., in the new technique described along FIG. 4-7 and 9, i.e., Steps 406 and 408 in FIG. 4, or elements 556, 558, 570 in FIG. 5. Purely by way of example and not limitation, consider in particular several core options to add a registration flow for imported instances to an MIaaS cloud:

First option: Run the provisioning flow for the closest related catalog image (which the analysis and plan phase 524 decided that the instance will be adjusted to simulate). Then swap the instance generated from the catalog image for the imported instance, on which most adjustments have been already done (e.g., additional patching, security compliance adaptations). This may be done on the image file level for suspended images. Some adjustments on the imported instance may be needed after the provisioning, e.g., if the provisioned instance had some identity (such as an IP address) that now needs to be transferred to the imported instance because the MIaaS management requires it but cannot tell it in advance. Note that such swapping and post-swap adjustment cannot be done by a user of an unmodified MIaaS cloud; it needs special authorization and, in order to be reasonably efficient and secure, automation tooling to be added to the MIaaS management system.

Second option (see FIG. 11): Modify the provisioning flow so that it still performs most of the registration work, but at some point picks up the image of the imported instance instead of the catalog image. If the provisioning flow performs real-time actions on the images (e.g., agent installs) those can also be reused (rather than being done in adjustments before the registration flow). If possible, retain the common parts of the provisioning flow and the new registration flow as actual common components (common subprograms, sub-workflows, or the like)

Referring now to FIG. 11, depicted therein is an exemplary provisioning flow in an MIaaS cloud environment (IBM SCE+ is one non-limiting example). In step 1101, a customer accesses customer portal and service catalog 1110, which are part of the BSS (568, 696) of the MIaaS. In one or more embodiments, “Adoption” or “import” is a new service type in the catalog 1110 (compared with the MIaaS cloud as designed without the new rapid migration mechanism) and is available for certain roles. In some cases, this aspect can be implemented by a relatively small change in an existing user interface (UI) of the MIaaS cloud. Appropriate parameters may be passed through to a provisioning process described below. The service request may involve, as shown, a service catalog; receipt of a request, fulfillment of the request, and updates to virtual machine provisioning status. In some cases, server resource usage status is shown during use.

In step 1102, provisioning is carried out with provisioning engine 1112. MIaaS Provisioning can include receiving the request from the portal 1110; using the image to create an instance; the creation of a virtual server including provision of resource requirements such as CPU, memory, disk, hostname, IP address, and/or sub-nets; installation of agents for management, configuration of agents, connection to virtual local area networks (VLANs); and/or server validation. One or more embodiments, with regard to registration of imported images as modification to such a provisioning flow, pick up the imported instance instead of a catalog image in the sub-step “Use image to create instance” of step 1102. Hence this sub-step becomes the creation of an instance 558 in the production zone from the imported instance, in relation to FIG. 5. If the landing zone is physically separated from the production zone, this involves actually copying the adjusted instance 554 from the landing zone to the production zone. If the landing zone is virtual, that instance may even be kept in place, but an address (name, pointer) of it is now given to the provisioning flow.

In one or more embodiments of the rapid migration method, during creation of the virtual server, the exact CPU, memory, and disk are created as in the imported instance or a plan from Step 524 (because the cloud may not have images with the exact resource sizes as the imported image, the next-larger sufficient size will typically be chosen). Some embodiments employ dual-homed instances which have a cloud services provider

internal IP address (for cloud management system of the MIaaS to access the instance) and a customer-facing customer address; this is typically different from the old IP address of the imported instance, and thus the address on that instance must be changed to the customer-facing address that the cloud management system provided.

5 It should be noted that words such as “must” or “shall” discuss items that are mandatory in an exemplary embodiment, but may be optional in other embodiments. No limitation to the claims is intended unless recited therein.

10 In some cases, with regard to the installation and/or configuration of agents, a new flow is provided for handling prior agents such as, for example, a performance monitoring agent or a backup agent, if the MIaaS cloud uses different agents or tools for the same purpose. Concrete, non-limiting, examples include Altiris products available from Symantec Corporation, Sunnyvale, CA, USA, and products available from NetIQ Corporation, Houston, TX, USA . The handling of such agents might include, for example, uninstall or an authorization reduction from administrator rights to user rights, 15 or a policy change in the tool. Alternatively, however, this flow can be done as adjustments in the landing zone before even starting the provisioning flow.

In step 1103, the virtual server is created. In the non-limiting example of FIG. 11, this can be carried out using the virtualization manager 1116 for INTEL technology (e.g. making use of VMware ESX VMs (VMware, Inc., Palo Alto, CA, USA) as shown at 20 1120) and/or the virtualization manager 1118 for POWER technology (e.g. making use of IBM PowerVM logical partitions (LPARs) as shown at 1120; International Business Machines Corporation, Armonk, NY, USA). In either case, suitable storage, such as storage area network (SAN) storage 1114 can be employed. Other types of technologies can be employed in other embodiments. This step can include taking the request from the 25 provisioning engine 1112; allocating appropriate computing resources such as CPU, RAM, disk space, and so on; copying the image over; placing same on the physical server; booting the operating system and middleware; and/or applying IP and hostname. In some embodiments of the adoption of such a flow for registration in rapid migration, a further step includes preparing custom access for tests. In MIaaS provisioning without 30 migration, such a test is not needed because this is a standard catalog instance, but here,

an imported, adjusted instance is being dealt with, and it will be appropriate to carry out the user acceptance test 560.

In step 1104, server validation is carried out with the aid of provisioning engine 1112. This can involve, for example, checking for correct installation and configuration of management agents; a security sweep (as to ports, password policies, and the like);
5 checking validity of user IDs; running a health check; documenting server information into a master database; and/or reporting to a suitable management tool such as the IBM Service Delivery Manager (SDM) available from International Business Machines Corporation, Armonk, NY, USA.

10 One or more embodiments advantageously employ automation for typical modifications in the adjustment workflow, so that these tests (e.g., security sweeps) typically do not fail – in fact these validation steps of the MlaaS cloud can be seen as the primary specification of what adjustments must be made in migration into this MlaaS cloud.

15 Step 1105 includes a final validation more on the BSS layer, in contrast to the OSS layer validation in Step 1104. This can include, for example, the SDM checking the validation report; confirming requests; approving release of the VM (i.e., instance) to the customer; the cloud services provider adding customer users to the servers (if not already there); and/or providing server access to the customer. In some embodiments, in
20 migration the access will be provided earlier for testing, as described for step 1103

Exemplary non-limiting details will now be provided regarding a number of steps that may be carried out in one or more embodiments.

Discovery (Step 518)

25 In some embodiments, with regard to a first run – before first transformation planning, standard tools such as IBM Tivoli Application Dependency Discovery Manager (TADDM) available from International Business Machines Corporation, Armonk, NY, USA, the IBM Galapagos tool, and the like, can be run well in advance of the actual migration on the production systems. With regard to the IBM Galapagos tool, see, e.g.,
30 Galapagos: model-driven discovery of end-to-end application-storage relationships in distributed systems, IBM Journal of Research and Development archive, Volume 52

Issue 4, July 2008, Pages 367-377 and Nikolai Joukov, Birgit Pfitzmann, HariGovind V. Ramasamy, Murthy V. Devarakonda: Application-Storage Discovery; SYSTOR 2010, Haifa, May 2010. In some cases, these tools can be preinstalled on an appliance (a small physical server) for easy installation into the customer environment 402. This approach
5 unifies these tools, customer inputs, and loading from certain available customer repositories. In at least some cases, the cloud services provider can ask about additional access to complete file system copies of instances, in particular virtual instances 522, already (but nothing in the process relies on that).

In some embodiments, with regard to a second run, the same is carried out in a per
10 wave fashion. That is to say, there is discovery again later, per wave, preceding migration by (for example) a couple of weeks. The same tools can be used, often with additional options and plug-ins to obtain more details. After this second run of discovery 518, in one or more non-limiting examples, a change freeze is implemented for code and configurations, so that the subsequent steps can rely that the information from discovery
15 is current.

Analysis and Planning (Step 524)

This phase decides whether business applications are capable of being migrated to one or more MIaaS cloud environments (SCE+ is a non-limiting example) and which
20 migration method to use. This is assumed to be integrated with general migration planning tools (for migrations and IT transformations of other types than to an MIaaS) that an organization might have, e.g., the IBM Migration Factory. Within this tool, the migration to MIaaS should share information such as for non-virtualizable software, software-OS-compatibility, and upgrade costs with non-cloud migration. Those might,
25 e.g., be available as static tables in a database underlying a planning appliance. One significant criterion for what workloads go into the path of rapid migration to an MIaaS cloud is that they already have the correct major OS version. For example, if the image catalog of an MIaaS cloud contained only images with Windows 2003, 2008, and Red Hat Enterprise Linux (RHEL) 5.6 as operating system versions, the planning might
30 decide to only use rapid migration to this MIaaS cloud for source instances 520 or 522 that already have one of these 3 operating systems, or at most additionally RHEL 5.1 to

5.5 (so that a minor upgrade is needed in adjustment). If there is non-virtualizable software on an instance that has one of these operating systems or other exclusion criteria (special hardware, potentially clusters), this instance still has to go to a different target environment, or be significantly remediated first (step 526). In some cases, regulated workloads are not migrated, but this is not intended to limit the scope of the claims unless expressly recited therein. In some cases, this aspect is implemented via extended static tables and discovery to minimize potential issues; although the skilled artisan, given the teachings herein, will appreciate that standard analysis techniques that analyze the possibility of physical-to-virtual techniques, suitably modified, can be employed to implement this aspect.

In one or more embodiments, this phase also includes analyzing compatibility of the patch status of the source instance with standards of the MIaaS cloud; this is specific to migration to MIaaS cloud.

In one or more embodiments, this phase also includes the selection of an image size. In at least some cases, simply retain the size, unless it is necessary to increase it to fit the new agents in that have to be installed during adjustments. If the MIaaS cloud only allows certain fixed image sizes (in terms of CPUs, memory, disk storage etc.), it may be necessary to resize the image later to the next larger suitable size; however, this should not be an exclusion criterion.

In one or more embodiments, this phase also includes wave planning, i.e., grouping servers that will be migrated together. This can be carried out, for example, using standard migration tooling (such aspects have been disclosed, e.g., in the aforementioned United States Patent Application Publications 20110213883 and 20120109844); using dependencies, locations, subnets, and so on; optionally, an attempt may be made at this stage to free hardware and the like. It should be noted that in the general case, there is no guarantee that there will not be cross-wave dependencies – some environments are too interlinked.

In at least some cases, security zones should be planned in this phase and certain cloud pools (i.e., groupings of instances that the cloud provides) and the like should also be planned. Refer to the discussion of the request GUI below with regard to FIG. 14.

Baseline Test and Backup (Step 528)

In the baseline test, the customer (or the migration team) performs all the tests that they will also apply after migration, to make sure the existing system passes all those tests. In one or more embodiments, the actual tests are the same as in user acceptance testing (UAT). This can involve, for example, some UAT tests that write into databases or the like such that they cannot be performed on production servers (e.g. production zone 574); test servers from the customer environment may then be used for the baseline test can then be employed, although the actual workload deployment there may be simpler than the production configuration.

In at least some cases, backup is performed by the customer, and the customer validates that it worked, as instance capture can cause problems. In some cases, the cloud services provider does not make backups on the customer side.

In some cases, the customer may carry out additional steps in this phase. For example, to enable the instance capture processes discussed just below, there are other things the customer should do, in particular:

- ensure there is a change freeze for their own developers;
- give the migration team (which can be in-house, from the cloud services provider, or another party) root rights on the instance to be migrated; and
- give the migration team a realistic change window for the migration, as physical-to-virtual (P2V) counts as a change, even in cases where it's run in the background.

Instance Capture

Multiple variants can be employed for this step. Refer to FIG. 12. Non-limiting examples will be presented, all of which are valid for arbitrary transfers to virtualized environments:

- If the source instance 520 in the customer environment 402 is still physical, a P2V operation (physical-to-virtual) will be involved as shown at 536. There are standard tools for this, e.g., the aforementioned PlateSpin tool, and the VMConverter tool available from the aforementioned VMware, Inc.

5 • If there are multiple drives, they may be treated by the same or different tools. A data drive transfer may, e.g., start earlier than a boot drive transfer and be done with a resynchronizing tool such as IBM Softek TDMF software, to accommodate large amounts of data with short change windows.

10 • If a P2V operation is needed, it might be local (i.e., first creating a virtual machine or image in the customer environment 402) followed by transport; or it might be integrated with the transfer, i.e., the new virtual machine or image file immediately exist in the target, i.e., the “landing zone” 572.

Significant customer-facing aspects in one or more embodiments include:

15

- The migration team needs to get root / admin rights on image
- If non-boot volumes are present, and transferred by a resynchronizing tool as described above, the tool must be installed beforehand.
- A change window should start here at the latest
- 20 • Applications are typically shut down during P2V time
- There should be a “starting zone” where the captured images go awaiting transport, in cases where all the transport is over a WAN 514 (at least for large images, small ones may go directly via the PlateSpin tool, for example).

25 Instance capture may include a suitable shut down process. In some embodiments, the enterprise applications (aka workloads) are shut down before a potential P2V step (if this has not already done during backup). This means to shut them down from the point of view of users – in a multi-component workload, it does not necessarily mean shutting all components down; e.g., just the frontend web server might
30 be disabled. It is safer not to even have any software components running during the capture step (cold P2V), but it is possible with the given tools to do it while things are
YOR920120450US2

running (warm P2V). The application owner, in one or more embodiments, defines what to shut down, and preferably carries out shut down, because the migration team, being tasked with infrastructure-level tasks, may not have full insight into this.

Instance capture may include, for example, the P2V step for physical servers. For already virtualized servers, in particular on VMware, little is required in this step. In one or more embodiments, after the end of P2V, changes are no longer allowed on the source (if this has not already been disallowed during backup), as they would not be transferred to the target. That is, typically the entire source instance is now shut down to ensure this. If the instance was already virtualized, it is shut down either before the image to be transported is finalized, or in an integrated transport, at the latest before the final resynchronization within that transport.

If a capture step is separate from the transport step, a “starting zone” should be provided for the captured instances, i.e., storage space for them that is close to the transport technology.

Transport to Cloud Services Provider (Including Data)

Data transfer may be over a network or via a physical medium, depending on network bandwidth and physical distance. This can be done, in some cases, with a resynchronizing tool (one example is IBM Softek’s TDMF), especially for large images or for non-boot drives. A tool such as TDMF should be installed beforehand, and if there is a large data set as non-boot drive, background data transfer may be started earlier than the P2V. If it is certain that the boot drive does not contain normal “data,” in at least some cases, it may be feasible to only do a change freeze on that, i.e., only resynchronize other drives. Due attention should of course be given to customer acceptability.

There should be sufficient bandwidth to the landing zone, typically much more than the standard WAN connection to the MIaaS cloud as it will be needed in the business-as-usual phase 711. In some case, special optical lines may be employed. It is possible to add WAN optimizers to same (such as those available from Riverbed Technology, San Francisco, CA, USA). In at least some embodiments, sizing should be for the maximum wave load, rather than average. Another option is to use physical storage media such as a high density SATA (Serial ATA (SATA or Serial AT

Attachment)) disk array and to transport it by courier or similar means. By using a resynchronizing tool to copy data onto them, resynchronization via the network is possible after that.

5 Landing Zone

One or more embodiments employ a “landing zone”, i.e., an environment in the same data center as the ultimate cloud place of the particular instance, and with the same hardware and hypervisor version (e.g., same ESX version for VMware ESX VMs), where such imported instances can be run. In some cases, with regard to the initial transport, if
10 it was done separately for an image file, all that is needed is enough storage space, as for any transport. However, at the latest for the first functional test, a runtime environment for the imported instance should be available. In one or more embodiments, there will be at least one landing zone per data center, as further WAN transfers are not employed once an instance is transferred to the cloud services provider. Potentially there may be even
15 more landing zones, in particular if the data center contains multiple physically separated or separately managed production zones; then there can be one landing zone per production zone.

Any security zoning of the customer that needs to be replicated can be done virtually, i.e., via the cloud service provider’s VLANs and firewalls, but unless the
20 customer requires physical separation of certain security zones also in the production environment (not usual in clouds), the instances can all land on the same hardware, controlled by the same management software.

As discussed before, the entire landing zone may be physically or virtually separated from the (corresponding) production zone. Exact placement and address are
25 pertinent considerations in one or more embodiments. In the case of virtual separation, one preferably imports each instance physically close to where it will stay after registration, so that no additional copying is needed. E.g., if the production zone is sub-structured into several units, where, e.g., each unit is controlled by a VMware vCenter Server(s) or a TSAM (IBM Tivoli® Service Automation Manager software available
30 from International Business Machines Corporation, Armonk, NY, USA), and the landing zone is virtual even within such units, the instance can be maintained on the same

physical server and storage when transferring it from the landing zone into the production zone. For example, the instances still belonging to the landing zone can have a flag “tolerated,” so that TSAM knows they are there but not yet managed in the production zone sense. However, if for additional security and lower complexity, the landing zone
5 resides in other units than the production zone, then actual copying is needed, and the registration flow, which is largely executed in the production zone, must be given an address of the instance in the landing zone and the authority to copy from such addresses (e.g., by opening a special port in a firewall).

10 Functional Testing

After transport to the cloud services provider, testing should be carried out to verify that the P2V and transport worked correctly. For this, standard methods employed in non-cloud P2V transformation can be used, as so far no MIaaS-specific adjustments have been made. A minimum test is that the instance can be started at all. In some cases,
15 this testing may be limited to OS, storage, and the like; in other cases, it may be extended to whether software or even the entire business application still works. In some cases, there may be a list of what was running from the time of shut-down. In some cases, a capability is provided by the migration team for automatic restart. On the other hand, in some cases, this is handled by the customer. The test is typically not identical to the
20 complete baseline test.

Backup / snapshot instance before/during adjustment

Directly after the transport into the cloud services provider, and optionally between any steps where tests are performed, a backup may be made (with a VMware
25 snapshotting technique) to enable easier remediation. However, in particular with mature processes, one may also decide not to make these intermediate backups in order to save time in the most common correct case. Backup is particularly useful before / after steps that involve a human decision, while automated steps can easily be re-done from the first backup.

30

Remediation 550

In at least some cases, this is a largely human step. If something goes wrong, it typically needs debugging. If too much difficulty is encountered, the instance and possibly the entire business application should be taken out of the current wave, i.e., the current weekend's (or other pertinent time period's) migrations. It is desirable to document such cases well, so that in the future such cases can be caught in the phases of Discovery 518 and Analysis & Planning 528.

10 Instance on Cloud Service Provider's VMware ESX virtual machines or the like

At this time, start the adjustment flow with the running instance that just passed functional testing.

Adoption and Adjustment Sub-flow 556

15 In at least some cases, adoption (referred to herein interchangeably as registration) and adjustment are carried by small modifications of the provisioning flow as discussed above in reference to FIG. 11. For a more general view, concentrating on the steps specific to migration (FIG. 11 showed the "non-migration" steps and the specific migration steps were discussed in the text) reference should be had to FIG. 13. In step 20 1302, determine whether any pre-adoption adjustments are needed or desired, i.e., adjustments that are done before and separate from the any registration steps, i.e., steps that involve the MIaaS Cloud BSS and OSS 568. This determining may refer as needed to a decision document or the like as output by the Analysis and Plan step 1306. In step 25 1304, start the adoption flow with appropriate instance parameters (and cloud zone) from analysis and planning step 1306. In step 1308, the portal flow passes the "adopt" parameter through, i.e., a parameter that indicates that the flow variant needs to be executed where an imported instance is registered. (This flow modification will be implemented in the system that is the basis of the MIaaS portal, for example, in IBM Tivoli Service Request Manager (TSRM) software available from International Business 30 Machines Corporation, Armonk, NY, USA). In step 1310, the provisioning takes place and the imported instance is picked up and any new adjustment sub-flows are called (can

be carried out, for example, in TSAM). In step 1312, a virtual server is created. (This is not needed if the instance remains the same, i.e., if the landing zone is purely virtual as described above.) (This can be carried out by the normal instance creation internal to the MIaaS cloud, for example, in TSAM). In step 1314, carry out server validation of the MIaaS cloud (assuming the MIaaS cloud has that and does not entirely rely on the prior steps to succeed without test); preferably with no changes, but more problems may be expected to be identified at this point for migrated instances than for instances provisioned from catalog images, and thus more time for remediation may be reserved.

In step 1316, prepare the instance for customer test access by providing appropriate keys to the customer. In step 1318, fulfillment of the request (i.e., the initial input to this flow) is validated if needed; this can be done via a shortened and automated process in one or more embodiments, where this is just a sub-flow in a larger migration flow, compared with the normal situation where this was a customer request to the MIaaS cloud portal.

The depicted sub-flow corresponds to step 556.

Backup / snapshots and tests (the neighboring overall steps) may occur between any of these sub-flows; even inside them, in particular in the provisioning flow 1310 after individual adjustments.

Start Adoption Flow with Instance Parameters 1304

In a non-limiting example, the portal is implemented in TSRM and includes a graphical user interface (GUI) and REST API (application program interface). In one or more embodiments, starting via REST API is believed to be appropriate as all parameters should be clear from the results of the analysis and plan step 524. Those may reside in a database or another well-defined repository so that they can be retrieved by software. While in FIG. 5, the analysis and planning component is depicted in the customer environment 402, where the discovered data arise, at some point the relevant part of the data for the adjustments step 556 should be transported to the MIaaS cloud environment 410, in particular its Cloud Management zone 680 (this arrow was omitted in FIG. 5 to avoid clutter). The transfer agent 682 and transfer core component 690 can take this on as

an additional task besides transporting images. However, in some cases, it may be desirable to initially start via a GUI in order to more easily track the activities.

In one or more embodiments, the results of the analysis and plan step 524 relevant for this sub-step should include the instance data, zoning plans, and the like so that the correct parameters can easily be input to the REST API or entered by a human operator into the GUI.

A naming / location scheme should be provided for instances in the landing zone so that at this point, information regarding which instance is to be registered can be input. In one or more embodiments, this is the name under which VMware vCenter Server knows the instance in question. In one or more embodiments, an automated check of the GUI inputs (size, OS) with the actual values of the instance is useful.

FIG. 14 shows an exemplary “Create Server” pop-up 1404 from an MlaaS portal GUI Portal Design self service center 1402. Migration servers are grouped into a project, and if each migration request can contain more than one instance (called image in this GUI), a “Configuration Set” option is provided. The user is given the option to fill in the name of the image he or she wants to use, or to click a magnifying glass to see a list of available images. VM size info should be part of the metadata for the image in one or more embodiments, possibly with an option to change same. The “migrate server” button starts the actual registration flow. The portal may now perform certain activities, such as recording the started flow in storage, and allowing its status to be queried from the GUI or the REST API.

Portal Flow - Passes “Adoption” Parameters Through 1308

Typically, the portal is a separate software component 1110 (in particular, part of the BSS) from the provisioning engine 1112, the component that carries out the actual registration (often a workflow engine, and part of the OSS), Hence the portal transfers the registration request to the provisioning engine. Compared with the normal provisioning request transferred between these two components, the choice of “registration” or “import” of a migrated instance, and the location / name of the instance is passed through, i.e., those are new parameters in that request in that API.

Provisioning Flow - Instance Pickup Step 1310

In one or more embodiments, based on a suitable naming / location scheme in the landing zone, and the provisioning engine 1112 (e.g., TSAM) obtaining the name / location of the instance to adopt from the portal 1110 (e.g., TSRM), the pickup is clear. In a non-limiting example, the name of the instance in an underlying VMware vCenter server may be used. This name might not be globally unique, but should be acceptable if the hostnames are used, which can otherwise be assumed to be all different already.

In embodiments where the landing zone is physically separated, certain firewall settings and a copying process to implement the actual pickup may be employed (which are a bit different from taking an image from the MIaaS image library).

Provisioning Flow – Adjustment Overview 1310

Turning now to FIG. 15, at 1502, note the customer instance after migration but before adjustment, i.e., the instance on a cloud hypervisor 554 in the landing zone directly after transport 538, and at 1504, the customer system after migration and adjustment (broken down into customer portion 1506 and cloud service provider portion 1508), i.e., as it will be in the production zone at 558. With regard to application/middleware layer 1510, little or no change is typically planned at this level (in particular because migration only to an MIaaS cloud, not a PaaS cloud, is considered in this example), except some configuration changes due to potential change of some basic OS configuration aspects such as IP address. In one or more embodiments, applications are not touched in this flow, inasmuch as P2V typically works well and this flow can be implemented with standard P2V techniques. With regard to application/middleware management layer 1512, little or no change is typically planned at this level, except some configuration changes in terms of security and due to potential change of some basic OS configuration aspects such as IP address.

With regard to OS management layer 1514, this flow includes removal of customer management suite components that are no longer required and policy adjustment of those that are; and installation of a standard management suite of the cloud

services provider as well as a cloud management tool of the cloud services provider (with regard to the latter, in at least some cases, agents are implemented in the previous aspect, and central tools are already present). In some cases, installation is implemented via a standard provisioning flow. Removal and policy adjustment are significant tasks.

5 It should be noted with regard to aspects 1512, 1514, that some tools may be common for the OS and App/MW layers.

With regard to OS layer 1516, in one or more embodiments, a minimum set of OS patches are installed as required by the cloud services provider to accept the Base OS (in an alternative embodiment, this can be implemented as a separate flow within otherwise
10 BAU cloud management). In some cases, an “open ports” check is carried out. In one or more embodiments, no change is needed at the base operating system level except for required OS configuration updates such as IP address.

Options to Handle Agents

15 In one or more embodiments, significant options, to be decided per agent type, include:

- Can certain agents be ignored? (Easiest flow)
- Must some agents be retained? In a non-limiting example, there may be a definite
20 sentiment to keep AD (Active Directory authentication product available from Microsoft Corporation, Redmond, WA, USA) and Altiris (for applications). In one or more embodiments, everything that is used for application management should stay, possibly with changes.
- Should some agents be deactivated? This means shutting them down, but not
25 uninstalling them. This is an option where agents no longer seem needed, but from the given information this is not quite clear yet and thus uninstalling is risky.
- Can / must some agents be completely uninstalled? In some cases, this may be appropriate to tidy images up, or because the agents need resources (e.g., NetIQ needs significant CPU for its monitoring). In one or more embodiments,
30 complete un-installation is never mandatory, as the customer can have all sorts of things on these servers and the cloud service provider does not judge that.

- Agents are typically not allowed to have root / admin rights. At this point, it should be determined whether those that must be retained will run without root / admin rights; and what more explicit rights do they may need. An example is Altiris to only patch applications.

5

In some embodiments, the cloud services provider may inquire of the customer as to how the customer employs these tools, what policies the customer has, where the tools are installed, and so on.

10 It should also be noted that any changes planned by the cloud services provider may need a significant new discovery effort – at least if the customer’s installs are as variable as those of other software, and as little standardized.

Minor Version, Fixpack, and Security Patch Updates

15 It is worth noting that the BigFix tool, now known as IBM Tivoli Endpoint Manager, available from International Business Machines Corporation, Armonk, NY, USA, is believed to be one particularly useful tool for security patches; however, other embodiments may use other tools.

Change of Addresses in Other Components

20 In a non-limiting example, the following address choices are made:

- The images retain their DNS names, i.e., during cut-over the customer will change the DNS entries for their domains. (This is much easier than changing the DNS names, which may be used inside many components of the application/middleware layer 1510 and its management 1512.)
- The instances get two new IP addresses, on new virtual NICs (network interface controllers), to enable management and backup tasks of the MIaaS cloud without interfering with the normal traffic over customer IP addresses and (virtual) NICs.
- The IP addresses from the customer environment are replaced by addresses in a subnet of the MIaaS cloud production zone; this is what should be entered in the above DNS entries.

30

Hence, in one or more embodiments, there is no need to change other servers or clients addressing the migrated instances, assuming the other servers or clients use DNS names (rather than using IP addresses directly); this has become increasingly common.

5 In some cases, where it is desired to reduce costs, the migration team cannot readily determine whether the customer used IP addresses for addressing. In such cases, it may be appropriate to assign to the customer a requirement to ensure that such is not the case. If this requirement is violated, problems may occur in user acceptance testing 560, and remediation will return the instances or entire workloads to the customer. This
10 problem is independent of the migration path or target, in particular of the fact that migration is to an MIaaS cloud.

Create Virtual Server 1312

In one or more embodiments, if it is possible to just retain the instance and at
15 most put it into a different VLAN via virtual switch and firewall changes, then it is not necessary to create any server at this point. Otherwise the virtual server creation works exactly as in standard provisioning of the MIaaS cloud.

Prepare Instance for Customer Test Access (Keys)

20 At this point in the adoption and adjustment flow, all significant things that need to be done on the instance itself have been carried out, and the customer is given access for testing the image as well. Most customer IDs should, at this stage, still be on the instance or the respective active directory (AD). Those that changed due to rights reduction of agents should typically be set already. In some cases, the cloud services
25 provider may need to hand out keys if access to the image is key-based. In one or more embodiments, this works the same as in the provisioning flow for catalog images described elsewhere herein.

Server Validation 1104

30 The validation as such can be unchanged from the provisioning flow for catalog images. In one or more embodiments, failure at this point is believed to be significantly

more likely in the case of a customer instance than in the case of a cloud services provider catalog image. In at least some cases, the cloud services provider may provide appropriate remediation.

It will be preferable to automate as much as possible of this validation and also of appropriate remediation if tests fail. Even better, however, wherever automated tests and remediation can be done at all, it may be advisable to undertake such aspects in the already in the adjustment step, even if they are repeated again in the server validation phase for compliance or audit reasons.

10 Request Validation 1105

In the migration scenario, where it is known in advance what will happen on each weekend (or other period of time), this aspect is preferably kept short and largely automated. The final approval from the customer after UAT may still need manual entry though.

15

At this point, the overall rapid migration flow from FIG. 5 can be resumed, after the details of adoption and adjustment (Step 556 and the subsequent transfer of the instance into the production zone, as shown in 558). Now the user acceptance test (UAT) 560 takes place (typically not only of one instance but of a business application supported by this instance and possibly others). If it fails, a manual decision should be taken, together with the customer, bearing in mind the following aspects:

20

- the issues may be minor ones that can be fixed locally in the cloud production zone 574 (e.g., DNS or firewalls not brought up correctly)
- it may be known that an adjustment failure is the issue - the particular adjustment may be skipped, e.g., a particular security patch. If the individual operation (such as an attempt to install a security patch) doesn't come with a safe roll-back, the cloud services provider can roll back to the previous stored instance. If there was a manual decision in this adjustment, the cloud services provider can, in some cases, try again with a different decision.

30

- In some cases, if it is not known that the issue is an adjustment failure, or issues cannot be repaired within the change window, or the variance goes beyond a certain threshold (e.g., more than five (for example) security patches can't be applied) then cancel the migration of this instance (and possibly of the entire business application). Restart the customer instances 520 or 522, typically based on the backup 532.

Cut-over and taking new instances into use (564)

If the UAT is passed, the instances of one business application or other migration unit in the cloud begin to be used by the customer. In the addressing choices are made as in the example above, the DNS of the customer is changed at this point; this can include changing reverse DNS entries – possibly as part of a standard procedure. In some cases, the change window can end here. If the customer had specific warnings, such as “application down” on a web entry page, they can now be taken off. Appropriate customer training can be provided if desired.

It is worth repeating that one or more embodiments offer the advantages of traceability and/or roll-back capability. Furthermore in this regard, traceability is a significant advantage and it is quite helpful for managed environments to have every step and especially adjustments done to the image clearly documented. The adjustment engine as well as TSAM/TSRM have the ability to create instance-specific logs for fully-transparent/change-controlled migrations. With regard to roll-back capability, the adjustment engine operates in a step-by-step mode and offers the ability to roll-back a certain adjustment if something went wrong. That ability together with the snapshots creates a relatively safe environment to execute the needed migration/adjustment steps.

Recapitulation

Given the discussion thus far, it will be appreciated that, in general terms, an exemplary method, according to an aspect of the invention, includes the step 404 of transferring a foreign instance 684 from a customer environment 402 to a target infrastructure-as-a-service cloud environment 410 as an image. This step can be carried

out, for example, via component 690, optionally by interfacing with agent 682. A further step 406 includes adjusting the foreign instance to standards of the infrastructure-as-a-service cloud environment to obtain an adjusted instance. This step can be carried out, for example, via component 692. A still further step 408 includes registering the adjusted
5 instance into a management system (generally 410; specifically 696, 698 in one or more embodiments) of the infrastructure-as-a-service cloud environment.

In general, the adjusting may include adjusting the foreign instance while it is running, or alternatively, adjusting the foreign instance as the aforementioned image.

In some cases, the registering includes running at least a portion of a provisioning
10 flow (e.g., as shown in FIG. 11) for a standard catalog image of the infrastructure-as-a-service cloud environment (the standard catalog image is selected for similarity to the foreign image); and swapping the image of the foreign instance for the standard catalog image. In general, such swapping may be carried out during the provisioning flow or after the provisioning flow; when carried out during the provisioning flow, it is referred
15 to as picking up the image of the foreign instance, during the provisioning flow, rather than swapping. Furthermore, swapping “after the provisioning flow” should be understood to include being after step 1102 (and optionally after most or all of step 1103, but before 1104).

Furthermore in this regard, consider three ways to proceed. Recall that
20 provisioning flow for catalog images is shown in FIG. 11. In some cases, proceed through the provisioning flow almost to the end like steps 1101-1103 to obtain a provisioned instance; then discard that provisioned instance and replace it with what is being migrated. (This is referred to as “swapping”) In some embodiments, the catalog image is never really put on the server. At the point where the catalog image would be
25 copied, pick up the instance and let the rest of the provisioning flow work on the migrated instance so there is less work to do at the end (FIG. 13 step 1310). In one or more embodiments, this point is in the sub-step “create virtual server” of step 1102. Such creation is detailed in step 1103. Within the virtual server creation, this point is in the sub-step “copy the image over” – in this aspect, the image of the imported instance
30 instead of the catalog image is what is copied.

As noted, in one or more embodiments, the infrastructure-as-a-service cloud environment is a managed-infrastructure-as-a-service cloud environment. Thus, in such cases, the transferring includes transferring to the managed-infrastructure-as-a-service cloud environment; the adjusting includes adjusting to standards of the managed-
5 infrastructure-as-a-service cloud environment; and the registering into the management system includes registering into a management system of the managed-infrastructure-as-a-service cloud environment.

As noted elsewhere, in the adjusting step, in some cases, the standards include one or more of security standards; infrastructure standards; patch managements standards;
10 and infrastructure management system standards (e.g., standards of tools such as 696, 698).

In some cases, an additional step 524 includes analyzing the foreign instance for suitability for transfer to the target infrastructure-as-a-service cloud environment; in such cases, the transferring is responsive to the analyzing yielding an affirmative result.

15 As noted, in some cases, a plurality of foreign instances are treated as a whole in the analyzing step.

In some cases, an additional step includes carrying out testing (refer, e.g., to decision blocks 404, 406, 408; in general, such testing can be carried out, for example, before the transport, after the transport, between different adjustments, after adjustments,
20 and after migration. A further step includes carrying out at least one of remediation and back-out if the testing fails (“Not OK” branches of blocks 404, 406, 408). As noted elsewhere, in some cases, a plurality of foreign instances are treated as a whole in the testing step.

In some cases, the testing includes a baseline test 528 prior to the transferring
25 step; and a user-acceptance test 560 after the registering step. The user acceptance test has at least one test case. The baseline test validates that the foreign instance initially passed the at least one test case of the user-acceptance test.

In another aspect, an apparatus includes a memory (e.g., RAM 30, cache 32); and at least one processor 16, coupled to the memory, and operative carry out or otherwise
30 facilitate any one, some, or all of the aforementioned method steps. Optionally, the apparatus further includes a plurality of distinct software modules 42. Each of the

distinct software modules is embodied on a computer-readable storage medium, and the distinct software modules include a transfer core component module, an adjustment component module, and a registration component module; as discussed elsewhere herein. The snapshot management system module optionally includes one, some, or all of the sub-components in FIG. 10.

In some cases, the transfer core component module obtains the foreign instance from a transfer agent in a customer environment and locates the foreign instance within a cloud landing zone of the infrastructure-as-a-service cloud environment. In some embodiments, the at least one processor is operative to instantiate the adjusted instance into a cloud production zone of the infrastructure-as-a-service cloud environment.

In another aspect, in some cases, the infrastructure-as-a-service cloud environment includes a cloud landing zone, a cloud production zone, and a cloud management zone. In at least some such cases, the at least one processor is operative to transfer the foreign instance into the cloud landing zone under control of the cloud management zone; the at least one processor is operative to carry out the adjusting on the foreign instance in the cloud landing zone under control of the cloud management zone; the at least one processor is further operative to instantiate the adjusted instance into the cloud production zone; and the management system forms at least a portion of the cloud management zone.

In some cases, the cloud production zone includes cloud production zone computer hardware of a first type and cloud production zone hypervisors of a first type running on said cloud production zone computer hardware of said first type; and the cloud landing zone includes cloud landing zone computer hardware of the first type and cloud landing zone hypervisors of the first type running on the cloud landing zone computer hardware of said first type. Fewer standards are imposed on instances managed within said cloud landing zone than on instances managed within said cloud production zone, and said cloud landing zone is separated from said cloud production zone via at least one of physical and logical techniques so as to avoid impact of said cloud landing zone on security of said cloud production zone. Stated in another way, the cloud landing zone is a set of computer hardware (e.g., servers, storage devices, etc.) and hypervisors of the same type as in the cloud production zone, but with fewer standards on the instances

managed there, and physically and/or logically separated from the cloud production zone so as not to impact the security of the cloud production zone.

MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

5 One or more embodiments advantageously provide techniques for management infrastructure analysis for cloud migration. In current migration analysis techniques, in particular the automated parts thereof, the focus is on key middleware such as application servers or databases. One significant advantage of the cloud computing paradigm lies in simplified management, which mainly comes from standardized infrastructure software
10 and standardized management processes. One or more embodiments advantageously extend automated migration analysis to these management aspects.

In an Infrastructure as a Service (IaaS) cloud, the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can
15 include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls). A Hardware-Infrastructure-as-a-Service (HIaaS) cloud provides bare-bones virtual machines as a service. It may also provide an operating system (OS) and even
20 software, but no support is typically provided for the OS or software.

On the other hand, a Managed-Infrastructure-as-a-Service (MIaaS) cloud provides full-service virtual machines. The service may, e.g., include OS patching and support for security and compliance of the OS. One significant aspect of MIaaS is simpler management via: standardization to a certain set of catalog images from which instances
25 are generated, automatic linkage of these instances to management tools during deployment, and/or not giving customers administration privileges on the OS level, so that the operating systems on the instances remain as the cloud administrators configure them.

Heretofore (prior to the emergence of the cloud computing paradigm), there was
30 much less need to analyze for compatibility with infrastructure standards, or for the costs associated with transformation thereto. One aspect of the cloud computing paradigm

involves enhanced efficiency through the standardization of infrastructure. The move to cloud computing also reflects a trend wherein information technology (IT) running costs have outpaced IT investment costs. To gain large benefits from investments in cloud computing resources, one or more embodiments advantageously facilitate migration of existing environments to such standardized infrastructure management systems and processes, starting with analysis.

Reference is made to the definitions of IaaS clouds, HIaaS cloud, MIaaS Clouds, Instance, Source instance, Image, Catalog image, Provisioning, Reinstall migration, Image-based migration, and Rapid migration to MIaaS cloud as defined elsewhere herein.

Refer now to FIG. 16, and note cloud description 16300; discovery and analysis of a source system to be migrated, shown at 16100 (an exemplary source system *per se* is shown in FIG. 17); and an infrastructure comparison engine 16200 to carry out comparison (referred to also as analyzing, mapping, and/or matching) between the two. Cloud description 16300 includes cloud infrastructure software standards and/or configurations (“configs”) 16310, 16320. Also included therein is an application-level description of cloud infrastructure management 16330; for example, a service-level agreement (SLA) offered by a cloud, or non-functional requirements that can be set as part of an application program interface (API). Optionally, description 16300 also includes include cloud infrastructure management processes 16340.

With regard to discovery and analysis 16100 on the source system to be migrated, step 16110 includes discovery of source infrastructure clients and servers; step 16120 includes discovery of source infrastructure configs and/or logs; and step 16130 includes obtaining (for example, by derivation) of an application-level description of the source infrastructure management. Optionally, process 16100 also includes step 16140, namely, discovery and/or learning with regard to infrastructure management processes of the source.

Process 16200 includes a comparison between the source and the target cloud. Step 16210 includes a matching between the source and cloud software, optionally taking transformation costs into account, and based on the cloud infrastructure software standards 16310 and the discovered source infrastructure clients and/or servers from step

16110. In one or more embodiments this matching process involves looking for conflicts as described elsewhere herein:

- a source infrastructure management component manages at least one object that at least one mandatory target infrastructure management component will manage in the target cloud infrastructure
- a source infrastructure management component uses at least one resource that at least one mandatory target infrastructure management component will use in the target cloud infrastructure
- current lack of a mandatory target infrastructure component on a client
- existence of a mandatory target infrastructure component in a different version
- existence of a mandatory target infrastructure component with different configurations.

In step 16220, map the infrastructure configurations based on the cloud infrastructure config standards 16320 and the discovered source infrastructure configs and/or logs from step 16120. In step 16230, map the application-level descriptions (e.g., descriptions of non-functional requirements as discussed elsewhere herein) of infrastructure management from the cloud (from 16330) and the source (from 16130). In optional step 16240, map the infrastructure management processes from the cloud (from 16340) and the source (from 16140).

Again, comparison and/or analysis 16200 steps can involve, in one or more embodiments, examining for conflicts.

With reference to steps 16340, 16140, and 16240, comparing the actual management processes is optional. In particular, if a radical migration such as to an existing cloud is made, then typically only the application-level descriptions need to be mapped. With regard to step 16120, in one or more embodiments, infrastructure configurations are discovered for mapping to potentially different software in the target cloud (purely by way of a non-limiting example, HP event filters associated with products from Hewlett-Packard Company, Palo Alto, CA, USA might be used in the YOR920120450US2

source environment and might be mapped to IBM Tivoli Monitoring (ITM) event filters, available from International Business Machines Corporation, Armonk, NY, USA, which might be used in the target environment). With regard to step 16130, several source aspects that are usually obtained by human interviews can be derived from the detailed
5 discovery, such as high-availability configurations, planned downtimes, and patch management setups. Even if still obtained by human interviews, they can be considered in automated analysis in step 16230 if the results are in a fixed format. With regard to step 16230, the application-level description of the current infrastructure typically has to be reached or exceeded, unless the business owners explicitly agree to a lower level.

10 In an alternative approach, instead of having a cloud description already given, derivation of the most suitable cloud description for the source environment may be part of the engagement, e.g., when the overall goal is to build a suitable private cloud for an enterprise.

By way of review and provision of additional detail, before performing any IT
15 migration, transformation, modernization etc., it is advisable to analyze compatibility between components to be retained and new components to be used. This may include, for example, selection of what to retain and what to change, as well as effort analysis of changes to be made to achieve compatibility.

Current migration analytics, in particular the parts done with well-defined
20 methods or even in an automated way, focus on compatibility on the operating system level, and sometimes on compatibility of key middleware such as application servers or databases. That is to say, they analyze whether an operating system version will still run on new hardware, or whether a database version will still run on a newer or related operating system version.

25 With the advent of clouds, in particular Managed-Infrastructure-as-a-Service (MIaaS) clouds, current migration analytics are insufficient, as these clouds also prescribe infrastructure management tooling and processes, as well as certain standards on servers that allow the infrastructure management to operate in a standard way.

One or more embodiments advantageously provide a system and method for
30 automated migration analysis for infrastructure management aspects, primarily for cloud migrations.

One or more embodiments advantageously provide a systematic and/or automated method for infrastructure management analysis.

Analyzing for infrastructure management migration is different from analyzing for middleware compatibility, because a significant question is whether the source
5 infrastructure management software conflicts with the standard infrastructure management software in the cloud (in particular MIaaS) and needs to be removed or modified (as opposed to whether the source infrastructure management software will run on the cloud operating system(s), which is what would be analyzed if the analysis were for middleware compatibility). Furthermore, other settings related to infrastructure
10 management standards need to be analyzed, such as assumed minimum patch levels.

In MIaaS cloud, migration implies the need for infrastructure analysis. One or more embodiments advantageously provide the ability to systematically analyze infrastructure changes needed and to apply them in a “rapid migration” process as described elsewhere herein.

15 Referring now to FIG. 19, in some instances, an exemplary method is implemented in two phases. In a first phase 1902, the to-be-analyzed elements, in particular the source infrastructure clients and servers from all source systems may be aggregated and discussed with the customer (i.e., the owner of or party responsible for the source system). For example, a cloud services provider may indicate to a customer
20 that it is necessary to uninstall all of the customer’s “Vendor-X-patching systems.” In a second phase 1904, the decisions are applied to individual source systems, and possibly refined if no general decision was possible in the first phase.

In one or more embodiments, it is believed advisable to detect conflicts pertaining to conflicting management in the first phase. Unless a comprehensive list of all potential
25 infrastructure software is available (this would be the case if Phase 0, numbered 1901, could already fill in the database 1908 very well, and all software in the customer’s source environment were already in such database), it is appropriate to analyze the overall software list at this time to see if there are such conflicts. Over time, the aforementioned database 1908 of conflicting and non-conflicting software can be
30 developed.

In some instances, the discovery of source infrastructure software 16110 can be done by carrying out discovery on the actual machines to be migrated, which contain clients of this source infrastructure software (e.g., asset management agents); by detecting servers of this source infrastructure software (e.g., a central asset management server);
5 and/or from questionnaires of deployed infrastructure tooling.

Thus, one or more embodiments advantageously provide a method for management infrastructure analysis for cloud migration including discovering infrastructure components in the source infrastructure; and analyzing the discovered infrastructure components for conflicts with mandatory infrastructure components in the
10 target infrastructure.

In some cases, the infrastructure components include one or more of infrastructure clients, infrastructure servers, infrastructure configurations, and infrastructure processes.

In some embodiments, conflicts include those where a source infrastructure component manages the same objects that a mandatory target infrastructure component
15 will manage.

In some cases, conflicts include those where a source infrastructure component uses the same resources, in particular a port, as a mandatory target infrastructure component uses.

Non-limiting examples of conflicts include one or more of the current lack of a
20 mandatory target infrastructure component on a client, the existence of a mandatory target infrastructure component in a different version, and existence of a mandatory target infrastructure component with possibly different configurations.

In some embodiments, the result of the analyzing step includes one or more of the recommendations to uninstall an infrastructure component, to install an infrastructure
25 component, to modify the configuration of an infrastructure component, to exclude servers with a particular component from migration, or to further discuss with the customer about an infrastructure component. Furthermore in this regard, the last two outcomes may occur in particular if a source infrastructure component is in a conflict, but likely to also be used to manage objects that are not under control of the cloud, e.g.,
30 middleware in a migration to an MIaaS cloud.

In some cases, the analyzing step is done in two phases; a first aggregated phase where all discovered infrastructure components are aggregated and generic recommendations given for them, and a second per-instance phase where the recommendations are applied to individual source systems, and refined if the general recommendation was not conclusive.

The analyzing step can, in at least some embodiments, include configuration comparison and/or process comparison.

With regard to configuration comparison step 16220, in some cases, it is advisable to compare configurations and make decisions for changes. This is particularly advisable if infrastructure software is used to manage both operating system (OS) aspects and application aspects. For instance, patching software on a source server may be set up to patch the OS as well as databases and web servers. If this source server is migrated to an MIaaS cloud, then the OS-level patching needs to be switched off in favor of the cloud patching software, while the database and web server patching needs to remain. On the other hand, in other cases, the source software may be uninstalled, but the server owners may wish to retain certain policies, e.g., backup frequencies or types of monitored events.

Consider now aspects of customer dashboards 1707 as shown in FIGS. 17 and 18. A related alignment question that can occur with MIaaS migration is integration into overall customer dashboards. For example, the customer may have a “CIO (chief information officer) dashboard” that gives an overview of the status of applications and servers. The customer may desire to retain such a dashboard after migrating some servers to MIaaS, both because the dashboard also shows applications, and/or because not all servers were migrated, but the customer would like a joint overview over migrated and non-migrated servers. In these cases, adapters from the reporting features of the MIaaS software into the customer dashboard should be provided. This aspect is shown in FIG. 18 by linking dashboard 1707 also to the target management program 2, numbered 1806, which is provided by the cloud.

Pertinent observations will now be provided regarding OS versions, patches and security. Another infrastructure aspect that an MIaaS cloud will typically prescribe are precise OS versions. For the OS, it is typically appropriate to decide what level of deviation is tolerated, and what may be automatically updated prior to migration. For

example, it might be decided to tolerate deviating minor versions (e.g., v5.2 instead of 5.3) or not, and there are often automatic updates tools among minor versions. It is also appropriate to consider non-version OS types, such as releases or editions.

5 Patches may need to be added in some cases. However, if source servers with back-level patching are found, it is not unlikely that there were application problems with some newer patches. Ideally, an effort should be made to try to find this out from the application owners beforehand. Patch updates should typically be followed by careful application tests.

10 With regard to other security settings, there are a number of good practices for security such as not enabling certain services and access rights to certain OS files. The source server should be brought into compliance with such requirements. Ideally, existing deviations from such best practices are not required by applications (e.g., applications should not require that an operating system administration password is only 6 characters long), but this should be discussed with application owners beforehand too.

15 In one or more embodiments, there are at least two significant ways of deciding what to change regarding operating system settings and the like, and in particular security-related settings:

- 20 • Compare with clean cloud images from a catalog. In some instances, this approach may be problematic, inasmuch as application-level and OS-level settings are not cleanly separated on typical operating systems. For example, if a source image has a service at a port where the catalog image does not, it is typically necessary to validate whether that is a useful application service.
- 25 • Compare with a security validation mechanism if the cloud has such a mechanism, e.g., a service activation check. Settings validated here typically must be achieved. However, such a test in a cloud might be incomplete because the cloud might rely on the cleanness of its catalog images.

30 Thus, in some cases, there can be a general cloud migration analysis phase before both of the above-mentioned phases, wherein decisions are taken as to which security settings are mandatory. Phase 0 in FIG. 19, numbered 1901, depicts a general cloud

migration analysis for management software; a similar scheme as depicted therein can also be used to carry out a general cloud migration analysis for OS settings.

Continued reference should be had to FIG. 16, and reference should also now be had to FIG. 17. FIG. 17 shows an exemplary source environment 1702 on which
5 discovery and analysis could be carried out. Environment 1702 includes a server 1703 with a management program 1, numbered 1701. Environment 1702 also includes a server 1704 with a management program 1706 interfacing with logs and/or configurations in a data store 1708. Physically, data store 1708 may or may not reside on the same virtual or physical server as the management program 1706. That is to say, the
10 configurations 1708 could also be inside server 1704; however, with a large management program, it is indeed possible that a configuration database 1708 is elsewhere (i.e., outside server 1704).

Two management programs are shown in FIG. 17 for brevity; in a typical scenario, long lists will be present. The management programs 1701, 1706 could
15 correspond, for example, to “backup management” and “compliance management.”

Environment 1702 also includes (in this example) two instances 1709, 1716. These are instances similar to those migrated in the rapid migration approach discussed elsewhere herein, i.e., they can be physical or virtual. The servers 1703, 1704 can in principle be virtual too, and a few of them might be migrated, so in one view they are
20 also instances. One or more embodiments, however, address the case wherein the cloud has its own management programs (in the cloud management zone 1899 of FIG. 18) so the servers 1703, 1704 are not migrated.

Environment 1702 includes one or more clients 1710, 1713, 1732 managed by management programs 1701, 1706. These clients may, for example, reside on instances
25 1709, 1716 and may be clients in the sense of how they relate to management programs 1701, 1706, while instances 1709, 1716 may also function as servers on which one or more applications 1712, 1714, 1718, 1720 run. Instance 1709 is managed by both management programs 1701, 1706, while instance 1716 is managed only by management program 1706. This demonstrates the desirability of per-instance analysis.

Instance 1709 has OS settings 1711 and instance 1716 has OS settings 1717. Client 1710 has configuration file 1715; client 1713 has configuration file 1731; and client 1732 has configuration file 1733.

In FIG. 18, which depicts an exemplary target (cloud) environment 1802, it is assumed that the management program 1, numbered 1701, was retained (and thus its server was migrated in the normal way into the cloud production zone - see instance 1803), while program 2 was exchanged for a corresponding program 1806 in the cloud (particularly, on server 1804, in the cloud management zone 1899). Program 1806 utilizes logs and configurations 1808, in a similar manner as program 1706 used logs and configurations 1708.

FIG. 18 also shows that OS settings changed somewhat (security, patch, or the like); this is illustrated by different numbers 1811, 1817 as compared to settings 1711, 1717 in FIG. 17. Instance 1809 is the analog to instance 1709 in FIG. 17, i.e., the migrated version of instance 1709; it includes unchanged client 1710 with config 1715 managed by retained management program 1, numbered 1701. It also includes client 2, numbered 1813, with config 1831, and managed by program 2 numbered 1806. Instance 1816 is the analog to instance 1716 in FIG. 17; it includes client 2, numbered 1832, with config 1833, and managed by program 2 numbered 1806. Applications 1712, 1714, 1718, 1720 now run in target environment 1802.

In some cases, instead of retaining program 1, or as program 3, it is also possible that the same management program already exists in the cloud management zone. Then the corresponding client can be retained, and the issue becomes whether the “config” also stays the same or is changed to cloud standards for such “configs.”

FIG. 19 shows exemplary phases in management infrastructure analysis for cloud migration. There is essentially a preparatory Phase 0, numbered 1901, followed by a two-phase analysis process including Phase 1, numbered 1902, and Phase 2, numbered 1904. In Phase 0, numbered 1901, at 1905, based solely on cloud infrastructure standards, make abstract decisions (for example, “no other OS patch management software may run”). At 1907, apply abstract decisions to some well-known infrastructure software packages. The result is (at least a partial) database 1908 of management clients

(versions) and general decisions for them; database 1908 can be updated later as explained below.

During Phase 1, numbered 1902, at 1911, discover source infrastructure clients and servers on each instance. In step 1912, aggregate the discovered versions. In step 1913, make a general plan per management client version (part of matching step 210, but not per source instance). Step 1913 takes input from step 1905, step 1912, and database 1908, and also updates database 1908 with its outputs. Step 1913 results in plan 1914; it will be understood that plan 1914 is illustrative and non-limiting.

During Phase 2, numbered 1904, per-instance decisions are made. In optional step 1916, if carrying out Phase 2 at a later time, re-discover infrastructure clients on the given instance (e.g., instance contains management clients 1, 2, 5 mentioned in exemplary plan 1914). If no re-discovery is done, such knowledge is still available from step 1911. In step 1917, apply the plan to these instances. In the non-limiting example, decide to uninstall client 1, retain client 2. Step 1917 thus takes input from optional step 1916 (else from Step 1911) and from plan 1914. In step 1918, where plan 1914 indicated “needs per-instance decision,” perform further analysis to decide (in the non-limiting example, decide for client 5). In optional step 1919, proceed to infrastructure configuration mapping (as in step 220). In the example, this is done, possibly, for clients 2 and 5.

FIG. 20 shows a possible user interface for a decision about a conflict, in particular for Phase 0 (Step 1907) or Phase 1 (Step 1913) where a general recommendation for a certain management software type is made. At 2002, a query is made whether the management software entered in the field “Middleware name” conflicts with any software in the cloud SCE+, version 1.5. The response is shown at 2004: in this case, there is indeed a conflict, and an explanation of the type of conflict is given. In another usage, the port that a certain piece of software is running on might be input (in the lowest input field in view 2002) to check for port conflicts.

The following code snippet in SQL is a simplified version of code that could be used for decisions in Phase 2 (Step 1917). Assume that software (management software and applications) installed on the source instances is gathered in a table “MIDDLEWARE_INSTALL_MASTER”. Furthermore, assume that component 1908 is

realized as a database table “MW_MANAGEMENT_LOOKUP,” where it is assumed for simplicity that everything with any conflict(s) is uninstalled, and precisely the software with known conflict(s) is collected in this table. The appropriate decision to make is whether to uninstall or not for each row in the former table, based upon whether the
5 middleware occurs in the second table. The following query does this, assuming the decision goes into column “FUNCTION” of the former table:

```
update MIDDLEWARE_INSTALL_MASTER set FUNCTION = 'UNINSTALL'
WHERE EXISTS
10 (SELECT 1
    FROM MW_MANAGEMENT_LOOKUP
    WHERE
    UPPER(MIDDLEWARE_INSTALL_MASTER.MW_DISTRIBUTION_NAME)
    = UPPER(MW_MANAGEMENT_LOOKUP.AGENT_NAME)
15 )
```

A possible result for a small subset of instances (given by column HOST_NAME) is shown in FIG. 21, by way of a non-limiting example.

Given the discussion thus far, it will be appreciated that, in general terms, an
20 exemplary method, according to an aspect of the invention, includes the step 16100 of discovering, in a source computing system 1702 having a source management infrastructure, at least one source infrastructure management component (examples given below). This step can be carried out, for example, using a suitable discovery tool, with appropriate modifications. In particular, known discovery tools cover infrastructure
25 software. However, where per-software modules are needed in the discovery software, the focus in existing packages may so far lie on middleware (databases, web servers, etc.) that new modules for infrastructure software should be added. Furthermore, it may be appropriate to extend the reporting features of the tools to include infrastructure software, and for specific reports that characterize infrastructure software so that step 16110 has an
30 output that makes steps 16210 and/or 16120 easy. Given the teachings herein, the skilled

artisan will be enabled to make suitable changes to existing discovery tool packages to implement one or more embodiments.

A further step includes obtaining a description 16300 of a target cloud infrastructure having a target management infrastructure (e.g. layer 64). The description includes at least one mandatory target infrastructure management component. This step could be carried out, for example, by using a suitable database program to query a description 16300 stored in a database; non-limiting examples include IBM DB2 software available from International Business Machines Corporation, Armonk, NY, USA, and ORACLE software available from Oracle Corporation, Santa Clara, CA, USA.

It may also be done by querying a cloud portal; if possible, via a programmatic interface, otherwise via one or more user interfaces. If none of these options is available it can be obtained manually from the cloud service provider. If that is also not possible, a further option includes provisioning an image from the cloud catalog, and performing discovery on that image to see what infrastructure software is installed on same. This infrastructure software is likely to be the cloud standard.

Thus far, it has been assumed that migration to a given cloud with given standards is desired. In some instances, there may still be a choice between multiple clouds; then multiple descriptions may be given and step 16210 may additionally look for the best-matching cloud. In another aspect, the goal may be to build a private cloud for the given source environment; then, a suitable cloud infrastructure software standard may be derived in step 16210 from analysis of the discovered source infrastructure. For instance, if it is discovered that 80% of the source servers are monitored with monitoring software X, and 20% with other monitoring products, then monitoring software X may be selected as the standard for the new private cloud, and monitoring software changes will only be needed on the remaining 20% of servers.

A still further step includes analyzing the at least one source infrastructure management component to determine whether at least one conflict exists with the at least one mandatory target infrastructure management component. This step could be carried out, for example, using infrastructure comparison engine 16200; refer to description of block 16210 elsewhere herein.

In the discovering step, the at least one source infrastructure management component could include, for example, at least one source infrastructure management client 1710, 1713, 1732; at least one source infrastructure management server (1704 with program 1706) (clients and servers can be discovered in step 16110); at least one source
5 infrastructure management configuration (e.g., in data stores 1708, 1715, 1731, and/or 1733); and/or at least one source infrastructure management log (e.g., in data store 1708) (configurations and logs can be discovered in step 16120). All the aforementioned elements things would typically be present, but they might not all be discovered. For example, the log might not be discovered. In at least some embodiments, when the focus
10 is on install / uninstall, it is sufficient to discover the client; and for the “Phase 1” analysis, it is sufficient to find the management programs.

In some cases, a further step 16130 includes obtaining a description of non-functional requirements of the source management infrastructure (e.g., with the aforementioned discovery tool) (also referred to herein as an application-level
15 description).

In some cases, in the obtaining step, the description of the target cloud infrastructure includes cloud infrastructure software standards 16310, cloud infrastructure software configurations 16320, and/or an application-level description 16330 of the target management infrastructure (i.e., a description of non-functional requirements on the
20 management infrastructure). These are typically given by the enterprise application supported by the currently considered instance. These can include, for example, an SLA or non-functional requirement, but SLAs can be seen as non-functional requirements as well. Other examples include change windows, up-to-dateness policies (such as for patches or in monitoring for problems), auditing requirements, and the like. Many of
25 these would typically be input manually. It should be ensured this is done in a fixed format to allow automatic comparison in step 16230. By way of a non-limiting example, if a cloud only provides certain patches every 4 weeks, but an enterprise application has a compliance requirement that requires such patches every week, then such enterprise application cannot be run on this particular cloud.

30 In cases as mentioned in the preceding paragraph, the analyzing step includes matching the cloud infrastructure software standards with the at least one source

infrastructure management client and the at least one source infrastructure management server, as at 16210; mapping the description of the non-functional requirements of the source management infrastructure with the description of the non-functional requirements of the target management infrastructure, as at 16230; and mapping the cloud
5 infrastructure software configurations with the at least one source infrastructure management configuration and the at least one source infrastructure management log, as at 16220.

As indicated by the arrows from block 16210 to block 16220 and from block 16230 to block 16220, in some cases, the mapping, at 16220, of the cloud infrastructure
10 software configurations with the at least one source infrastructure management configuration and the at least one source infrastructure management log is at least partially based on the matching, at 16210, of the cloud infrastructure software standards with the at least one source infrastructure management client and the at least one source infrastructure management server, and the mapping, at 16230, of the description of the
15 non-functional requirements of the source management infrastructure with the description of the non-functional requirements of the target management infrastructure.

In some embodiments, in the discovering step, the at least one source infrastructure management component further includes at least one source infrastructure management process, as shown at 16140; in the obtaining step, the description of the
20 target cloud infrastructure further includes at least one target infrastructure management process 16340; and, as seen at 16240, the analyzing further includes mapping the at least one source infrastructure management process with the at least one target infrastructure management process.

As indicated by the arrow from block 16230 to block 16240, in some cases, the
25 mapping of the at least one source infrastructure management process with the at least one target infrastructure management process, at 16240, is at least partially based on the mapping of the description of the non-functional requirements of the source management infrastructure with the description of the non-functional requirements of the target management infrastructure, at 16230.

30 In some cases, in the analyzing step, the at least one conflict includes the at least one source infrastructure management component managing at least one object that the at

least one mandatory target infrastructure management component will manage in the target cloud infrastructure.

In some embodiments, in the analyzing step, the at least one conflict includes the at least one source infrastructure management component using at least one resource (e.g. port) that the at least one mandatory target infrastructure management component will use in the target cloud infrastructure.

In some cases, in the analyzing step, the at least one conflict includes one or more of the current lack of a mandatory target infrastructure component on a client, the existence of a mandatory target infrastructure component in a different version, and the existence of a mandatory target infrastructure component with possibly different configurations.

Consider the current lack of a mandatory target infrastructure component on a client. For example, if the cloud requires event monitoring agent “ABC” in version 6.1 on all instances (i.e., ABC is a mandatory component like 1813), and a source instance 1709 does not contain ABC in any version, a decision would be to install ABC. Consider the existence of a mandatory target infrastructure component in a different version. For example, if source instance 1709 already contains ABC in version 4.5 or 6.0, but the cloud requires 6.1, then it may be necessary to upgrade (probably easy if it was 6.0) or uninstall and reinstall (maybe needed if it was 4.5, i.e., rather old).

Consider further the existence of a mandatory target infrastructure component with possibly different configurations. If ABC in version 6.1 was already there, or 6.0 was present and it has an upgrade feature to 6.1, it may have a local configuration specifying what events it monitors. On an instance provisioned from a cloud catalog image, a different configuration for ABC 6.1 may be present that monitors a different set of events. One solution is to overwrite the source configuration with the cloud standard configuration. Another is, in this example, to join the two configurations such that the union of the two sets of events gets monitored. This is possible if the corresponding central cloud management program, say “ABCServer” (corresponding to 1706), tolerates when unexpected events come in. If ABCServer is the only tool that processes these events, it may still be useless, but if it forwards events to a customer dashboard 1707, it might pass these unexpected events through. This could help if, e.g., the event monitoring

system(s) was used both to monitor events from an application 1712 (say crashes and restarts of this application) as well as OS-level events. The application 1712 remains under control of the customer and the customer might still want to see these events in the dashboard 1707.

5 In some cases, the analyzing step includes one or more of (i) recommending un-installation of a source infrastructure component; (ii) recommending installation of a target infrastructure component; (iii) recommending modifying configuration of an infrastructure component; (iv) recommending excluding servers in the source environment with a given source infrastructure component from migration; and/or (v) further discussing a given infrastructure component with a customer. The last two outcomes (exclusions (iv)/further discussions (v)) may occur in particular if a source infrastructure component is in a conflict, but the same source infrastructure component is likely to be used to manage objects that are not under control of the cloud, e.g., middleware in a migration to an MIaaS cloud (corresponding to application 1712, 1714 etc.). With regard to possibilities (i)-(iii), the infrastructure component(s) mentioned could be in the source environment prior to migration, could be adjusted during rapid migration and/or after transport, and so on. With regard to possibility (iii), the same might start on either side – typically modification takes place if there was both a source configuration and a standard target configuration for the same software, and some compromise is appropriate. With regard to (i), it could be decided that event monitoring software other than ABC is not desired and any such software other than ABC could be uninstalled.

It is worth noting that “further discussion” option (v) may involve human agency. However, the automatic engine may need to output this as a recommendation. For example, if in 90% of the instances it can make an automated decision, but in 10% of the cases it cannot, it is already well worthwhile using such an engine. But the engine cannot simply guess in the remaining 10% cases, it must admit that it doesn’t know. In some embodiments, these questions pop up in an interface with a structured format to answer them so that automation can proceed from there on. If several people need to decide a collaboration tool can be provided.

In some cases, the discovering includes discovering, in the source computing system, a plurality of source infrastructure management components; and the analyzing includes: initially aggregating all of the plurality of source infrastructure management components; developing generic recommendations for the aggregated plurality of source infrastructure management components; applying the recommendations on a per-instance basis; and refining at least some of the applied recommendations in cases where the generic recommendations are inconclusive. See, e.g., FIG. 19 and accompanying text.

In some instances, the discovering is carried out on at least one physical machine 1709, 1716 in the source computing system, and on this at least one physical machine is software 1710; 1713 or 1732 that is a client of the at least one source infrastructure management component. On the other hand, in some cases, the discovering is carried out on at least one physical machine in the source computing system, and the at least one physical machine is a server 1704 on which the at least one source infrastructure management component runs. Questionnaires can also be used.

In some cases, an additional step includes providing a system, including distinct software modules; each of the distinct software modules is embodied on a non-transitory computer-readable storage medium. The distinct software modules include a discovery tool module to carry out discovery 16100, a description module to hold and access description 16300 (see discussion above of obtaining description 16300), and an infrastructure comparison engine module to implement engine 16200 and carry out analysis. Examples of a description module include a suitable database program to query a description 16300 stored in a database; cloud portal query code (querying, if possible, via a programmatic interface, otherwise via one or more user interfaces); or code to provision an image from the cloud catalog, and perform discovery on that image to see what infrastructure software is installed on same. This infrastructure software is likely to be the cloud standard.

Non-limiting examples of discovery tools include IBM Tivoli Application Dependency Discovery Manager (TADDM) available from International Business Machines Corporation, Armonk, NY, USA, the IBM Galapagos tool, and the like. With regard to the IBM Galapagos tool, see, e.g., Galapagos: model-driven discovery of end-to-end application-storage relationships in distributed systems, IBM Journal of Research

and Development archive, Volume 52 Issue 4, July 2008, Pages 367-377. In some instances, with respect to the discovery tool at least a portion (e.g., scripts or agents) typically runs in the source environment 1702, while another portion runs on a dedicated server in the target environment. However, the entire discovery tool could run on a
5 dedicated server in the target environment with a connection or connections to the source environment 1702. In some instances, the entire discovery tool could run on the environment 1702. Engine 16200 and the database with description 16300 would typically reside in the target environment (e.g., FIG. 3). They can also reside in a separate analysis environment, e.g., on servers of the team that performs the analysis. As
10 discussed above, existing tools can be modified to implement aspects of the invention.

USE OF SNAPSHOTS TO REDUCE RISK IN MIGRATION TO A STANDARD VIRTUALIZED ENVIRONMENT

A hardware infrastructure as a service cloud (HIaaS Cloud) provides bare bone
15 virtual machines as a service. It may also provide an operating system and applications; however, support is typically not provided for any software installed on the virtual machine (VM). A non-limiting example is the Amazon Elastic Compute Cloud (Amazon EC2) available from Amazon Web Services LLC Seattle, Washington, USA. A managed infrastructure as a service cloud (MIaaS Cloud) provides full service virtual machines.
20 The service contract includes providing software services for the underlying virtual machine. These services include software upgrades, patch support, support for compliance, license management, and the like for the provided software infrastructure. A non-limiting example is the IBM SmartCloud Enterprise+ also known as IBM SCE+, cloud environment available from International Business Machines Corporation,
25 Armonk, NY, USA. Image-based migration uses customer images for migration of the application to the cloud. This approach minimizes expensive application re-installation. A cloud standard is the minimum set of requirements needed to be met by a customer image to be run and managed in cloud.

From an information technology (IT) management perspective, a hosting
30 environment (a cloud is a non-limiting example) that aggressively enforces standardization (e.g., at operating system (OS), middleware, and/or application layer) can

lead to significant management cost savings. However, migration of custom systems (called source systems) into such a standardized environment requires changes to be made to the software stack of the source system in order to comply with the standard. Those changes potentially carry a high risk of breaking the source system, rendering it unusable.

In the recent years companies have increasingly adopted virtualization to achieve hardware consolidation. One or more embodiments address migration into a standardized virtualized environment. One example of such a standardized virtualized environment is an Infrastructure as a Service (IaaS) Cloud. In one or more embodiments herein, a cloud environment is employed as a non-limiting example of a standardized virtualized environment. However, it is to be understood that one or more embodiments are broadly applicable to a general standardized virtualized environment that may or may not expose the same kinds of APIs of an IaaS cloud.

One or more embodiments advantageously provide a system that reduces the risk of migration via the use of snapshots to systematically capture system states at critical points of time during migration. Moreover, one or more instances provide a system that supports two types of snapshots: short term snapshots (STS) and long term snapshots (LTS), each optimized for specific usage scenarios. Furthermore, one or more embodiments provide a method that balances various requirements of “snapshotting” (resource utilization, performance impact, long-term preservation, and the like) through a combination of short term snapshots and long term snapshots. Even further, one or more instances provide a system that supports roll-back to any type of snapshot. Yet further, one or more embodiments provide a system that automatically provides an audit of the exact changes applied to the source system through the migration process. The changes can be viewed at two granularities: file level and product level. Two levels of information are available for a long term snapshot: the file system (including meta-data information about each file such as file name, owner, permission, modified date, and optionally a hash of the content of the file); and the products (software packages) contained within the system.

For the avoidance of doubt, “snapshots” as used herein include both long-term snapshots and short-term snapshots; however, the process of “snapshotting” refers to the

creation of a short-term snapshot while the creation of a long-term snapshot is referred to as checking in or check in.

The following is a non-limiting example of file-level and product-level changes referred to as “diffs”):

5

- <removed-file name="/Program Files/Altiris/Altiris Agent/Software Delivery/pkgdlv1k.tmp" digest="sha1:da39a3ee5e6b4b0d3255bfef95601890afd80709"/>
- <added-file name="/cygwin/usr/share/groff/1.19.2/tmac/mm/locale" digest="sha1:da39a3ee5e6b4b0d3255bfef95601890afd80709"/>
- <added-file name="/IBM/ITM/InstallITM/MiaError.log" digest="sha1:da39a3ee5e6b4b0d3255bfef95601890afd80709"/>
- <modified-file name="/Documents and Settings/Administrator/NTUSER.DAT" old-digest="sha1:33aca147bd43c253fbf904fea76aa63c1aab3b32" new-digest="sha1:7a12613062e36b27a4c20147ebd702671e0272ca"/>
- <added-product vendor="IBM" name="IBM ITCAM for Microsoft Applications Entry" version="6.2"/>

]

20

Advantageously, one or more embodiments reduce the risk of migration by providing a “time capsule” of the system being migrated, at critical points of time during migration, such that the system can be rolled back to a “known good state” in case of failures. A customer will also know what changes have been applied to its system during migration before the system is on-boarded into the cloud. The reduced risk translates to improved engagement success rate, reduced downtime of applications, and lower migration cost due to reduced debugging time during failures.

25

30

It is worth noting that virtual machine (VM) snapshot and roll back is a general mechanism available in modern hypervisors, which have been used in various scenarios, such as system debugging and/or system security. In one or more embodiments, the “short term snapshot” is essentially the same as a VM snapshot; however, the usage scenario is different (reducing risk in migration). Additionally, one or more

embodiments employ novel long term snapshots and the novel combination of the two types of snapshots.

It is also worth noting that traditional backup tools are used to back up a system and later on restore the system to the state before the back-up. In one or more
5 embodiments, the long-term snapshot is stored as an image object in an image library; existing backup systems deal with the concept of a file system, whereas an image is a super-set of a file system. An image also includes aspects of a virtual machine, such as central processing unit (CPU), memory configuration, and so on. An image library provides versioning and analytics capabilities.

10 Referring now to FIG. 22, one or more embodiments advantageously reduce migration risk by using an image library 2202. Here the “snapshot” action produces short term snapshots, so 2216-2228 are short term snapshots. The act of “check in” into the image library 2202 produces long term snapshots 2206, 2208, 2210. In one or more instances, the states at critical points of time are captured to form an image version chain
15 (e.g., version zero 2204, version one 2206, version two 2208, and version “n” 2210). In the case of a system malfunction, one or more embodiments revert to the “last known good state.” One or more instances retroactively “introspect” image versions (i.e., inspect the content of the system at a point of time in the past). One or more instances “diff” (i.e., carry out a comparison to identify differences) any pair of versions at the file level
20 to understand changes made between them. Furthermore in this regard, one or more embodiments, as part of the introspection aspect, reconstruct an earlier state. For example, when an image is checked into the image library, the contents (meta-information) are parsed and stored in the image library and the resulting data structure can subsequently be inspected. In the short term case (short-term snapshot such as in
25 VMware) a rollback process can be carried out.

Still referring to FIG. 22, the process begins with an original image 2212 (after physical-to-virtual (“P2V”) translation). This original image is imported into the image library 2202 as version zero, at 2204. An instance of the original image is started at 2230. A first adjustment is made at 2246, to transition from instance 2230 to instance
30 2234. Snapshots 2216, 2218 are taken of each incremental change; i.e., of instance 2232 and instance 2234 (instance 2230 is simply an instantiation of the original image after P-

2-V so no snapshot is needed, although one could be taken if desired). The snapshots are of a short term nature. Upon completion of the first adjustment 2246, snapshot 2218 is checked into the (long term) image library as version one at 2206. A second adjustment is made at 2248, to transition from instance 2234 to instance 2240. Snapshots 2220, 2222, 2224 are taken of each incremental change, i.e., of instances 2236, 2238, 2240. Upon completion of the second adjustment 2248, snapshot 2224 is checked into the image library as version two at 2208. An even further adjustment is made at 2250, to transition from instance 2240 to instance 2244. Snapshots 2226, 2228 are taken of each incremental change, i.e., of instances 2242, 2244. Upon completion of the third adjustment 2250, snapshot 2228 is checked into the image library as version “n” at 2210. As indicated by the ellipsis, any desired number of adjustments can be made. Time increases along the axis from left to right. Again, the arrows from 2232-2244 to 2216-2228 represent a short-term snapshotting process and the arrows from 2218, 2224, 2228 to 2206, 2208, 2210 represent a long-term check-in process.

Elements 2230-2244 represent the same running instance at different points of time. Each small change is snapshotted as at 2216-2228; however, only the completion of significant adjustments 2246, 2248, 2250 triggers the check-in of long-term snapshots 2206, 2208, 2210. The frequency of long-term and short-term snapshots is thus different.

Short-term snapshots are relatively short-lived (such as the duration of the migration window, which typically ranges from several hours to a weekend) and contain memory states. Short-term snapshots are employed for temporary debugging purposes, and are typically intended to be discarded after debugging completes. Short-term snapshots occur more frequently, thus requiring fast (on the order of seconds) operation. Short-term snapshots can be supported by a modern hypervisor; VMware is a non-limiting example (VMware software is available from VMware, Inc., Palo Alto, California, USA). However, each snapshot creates a new node in a snapshot chain. In almost all implementations of snapshots in modern hypervisors, the technique of copy-on-write is used where a copy-on-write child is created on top of a read-only parent, allowing a very fast creation operation (no data copying). Subsequent write operations only go to the child, whereas read operations go to both the child and the parent. When multiple snapshots are created for the same VM, a snapshot chain is formed where each

snapshot is a node in the chain. The performance of a read operation then depends on the length of the chain, because the hypervisor may have to traverse the chain all the way to the head, in order to retrieve the data being asked for.

5 In at least some instances, long chains lead to performance degradation and instability, and thus may not be appropriate for production systems. In one or more embodiments, short term snapshots 2216-2228 of virtual machine instances 2232-2244 are taken frequently and stored in the hypervisor as virtual machine images. Long term snapshots 2204-2210 are taken less frequently (e.g., at an initial state 2212 and after adjustments 2246, 2248, 2250) and checked in as images (2204), 2206, 2208, 2210 in
10 image library 2202. Long term snapshots resemble long term backups but are different in several ways.

In particular, with regard to long-term snapshots, in some instances, problems may occur long after migration (e.g., in steady-state); in such cases it is still necessary to revert to the “last known good state.” The customer may require a snapshot of the
15 original image before adjustment (e.g. for audit and/or compliance purposes). Use of an image library 2202 is more suitable in such circumstances. One or more embodiments thus store images in an image library as persistent data for long term preservation. In addition, the image library provides versioning and analytics capabilities for the stored long-term snapshots. However, long-term snapshots in an image library are typically less
20 efficient in operation (on the order of tens of minutes).

Referring now to FIGS. 23 and 24, in one or more instances, a snapshot management system (labeled in the figures as a “snapper” for brevity) 2340 manages both short-term snapshots 2346 and long-term snapshots 2344 (the snapshot management system also initiates rollbacks as described below). In one or more non-limiting
25 exemplary embodiments, short-term snapshots 2346 are VMware snapshots (VM images of VM instances) that reside on an operational repository 2468 in VMware data center 2342, while long-term snapshots 2344 are checked into the reference repository 2456 of virtual image library 2202. As used herein, an operational repository is defined as a storage location where a hypervisor stores image files corresponding to running,
30 suspended, and/or stopped VM instances. Snapshot management system 2340 creates and deletes short-term snapshots in the form of VM instances 2470 within VMware data

center 2342. Images of the instances are checked in and out of the operational repository 2468 as needed. The VMware data center is used for short-term snapshots for rollback in the production environment 2452. When a long-term snapshot is desired, an image in the operational repository is checked into the reference repository 2456 within analytics
5 environment 2450. Indexer 2458 indexes images in reference repository 2456, and optionally also indexes images in operational repository 2468 and/or VM instances 2470, to create knowledge base 2460. This knowledge base is used by engine 2462 to carry out analysis.

Furthermore in this regard, snapshot management system 2340 creates
10 (frequently) short-term snapshots of executing virtual machine instances 2470 (snapshot management system 2340 also deletes the short term snapshots from time to time). The short term snapshots are stored as virtual machine images of the virtual machine instances 2470 executing on the hypervisor, in repository 2468. When it is desired to create a long-term snapshot (e.g., prior to an adjustment), snapshot management system
15 2340 causes a corresponding one of the virtual machine images in repository 2468 to be placed in reference repository 2456. On the other hand, as seen in FIG. 25, when a rollback is desired, snapshot management system 2340 checks the image out of reference repository 2456 as a template, and places it into operational repository 2468 (as seen at 2572) and then reconfigures the template to match an existing VM instance and
20 instantiates a new virtual machine 2573 based on the template, which VM image 2573 is swapped with the existing VM instance 2574 for which rollback is desired.

For the avoidance of doubt, “check-in” refers to placing something into the virtual image library and “check-out” refers to taking something out of the virtual image library (with version control); these terms are not used in reference to the VMware data store
25 because in this context VMware does not have version control and what occurs is simply a file copy operation. In one or more embodiments, both long-term and short-term snapshots are triggered by adjustments. Short term snapshots are taken at each incremental step in an adjustment, but long term snapshots are only taken at the beginning and end of a significant adjustment process, skipping all incremental
30 adjustments in the middle. This is shown in FIG. 22. Other approaches could be used in other embodiments.

Note that the VIL 2202 in FIG. 24 is the same as the Image Library 2202 in FIG. 23; and the operational repository 2468 in FIGS. 23 and 24 could be, for example, a VMware data store.

FIG. 25 thus shows an exemplary rollback to a particular long term snapshot. In step 1, image 2576 is located in image library 2202 based on its long term snapshot identifier (LTS ID). In step 2, the image is checked out as a template 2572 in operational repository 2468. In step 3, the template is reconfigured to match the existing VM instance 2574, and a new VM 2573 is instantiated. In step 4, the new VM 2573 is swapped for the existing virtual machine instance 2574.

Again, to summarize, in one or more embodiments, rolling back to a long term snapshot involves the following steps:

1. Locating the image in the library given a long term snapshot ID
2. Check out the image from the library into an operational repository as an image template
3. Reconfigure template to match existing VM instance and instantiate a new VM
4. Swap the image file associated with the newly created VM instance with that of the existing VM

FIG. 26 shows an exemplary flow chart for a method to compute exact changes made to an image during migration. In step 2602, locate the long term snapshot of the original customer image (e.g. 2204). In step 2604, locate the long term snapshot after all adjustments are complete (e.g. 2210). In step 2606, compute the difference (see example above) between the two versions at the file and/or product level; then, in step 2608, return the results.

FIG. 27 shows an exemplary process flow. As seen at 2702, a snapshot is taken of an initial state, with a selective check-in (the asterisk next to a snapshot or check-in in FIG. 27 denotes a selective snapshot or a selective check-in as the case may be). Again, a snapshot creates short term snapshots, whereas check-in creates long term snapshots. In 2704, a first adjustment is made (e.g., adjustment 2222 in FIG. 22). In decision block 2706, a determination is made (manually and/or computer-aided) as to whether the first

adjustment was successful. If not, a selective snapshot and selective check-in are made in step 2708, and a rollback to the previous configuration is carried out in step 2710, with any required manual remediation carried out in step 2712. The first adjustment can then be re-tried if desired. On the other hand, if the first adjustment was successful, as seen at 5 2714, a snapshot is taken of the updated state, with a selective check-in. A second adjustment is then made in step 2716 (e.g., adjustment 2226 in FIG. 22). In decision block 2718, a determination is made as to whether the second adjustment was successful. If not, a selective snapshot and selective check-in are made in step 2720, and a rollback to the previous configuration is carried out in step 2722, with any required manual 10 remediation carried out in step 2724. The second adjustment can then be re-tried if desired. On the other hand, if the second adjustment was successful, the process can continue; for example, with a snapshot of the further updated state and a selective check-in, followed by any further adjustments (not shown).

One or more embodiments thus advantageously provide a system that reduces risk 15 in migration by using a combination of short term and long term snapshots to systematically capture system states and support roll-back to any snapshot. One or more instances provide a method for initiating short term and long term snapshots and separating them, as well as Virtual Image Library-based (VIL-based) long term snapshots. As seen in FIG. 25, one or more embodiments provide a method for rolling 20 back to a long term snapshot. One or more embodiments permit comparison of long term snapshots at the product and/or file levels. Further, as seen in FIG. 26, one or more instances provide a method for computing the product-level and file-level changes made to the custom image during the migration process, the output of which could serve as an audit record for clients that desire same. Even further, as seen in FIG. 27, one or more 25 embodiments provide a process that integrates snapshot and rollback with each risky operation (e.g., adjustments 2222, 2226, 2230) on a VM.

It is worth noting that, while the idea of a checkpoint has been around for a long time, it is mainly used to periodically save the states of a long running system so that if the system fails it can resume computation from the last checkpoint rather than starting 30 from time zero (as opposed to being used for cloud migration, as in one or more embodiments).

Given the discussion thus far, it will be appreciated that, in general terms, an exemplary method, according to an aspect of the invention, includes the step of, during a process of migrating a source system into a standardized virtual environment, snapshotting virtual machine instances 2470 of the source system executing in a
5 hypervisor 2342 as virtual machine images in an operational repository 2468 of the hypervisor. The virtual machine images in the operational repository are short-term snapshots 2216-2228. This step can be carried out, for example, by the snapshot management system 2340 invoking or issuing a command to the hypervisor (VMware is a non-limiting example of the latter). The method also includes, from time to time during
10 the migration process, creating long-term snapshots 2344 of the source system. The long-term snapshots are created by checking given ones of the virtual machine images from the hypervisor operational repository into an image library 2202 as image objects (e.g., in reference repository 2456). This step can be carried out, for example, by the snapshot management system 2340 invoking or issuing a command to the virtual image
15 library 2202.

Furthermore with regard to the step of snapshotting, snapshotting mechanisms *per se* are typically provided by known hypervisors. However, in one or more embodiments, snapshot management system 2340 issues a command to or invokes the hypervisor. That is to say, existing hypervisors have snapshotting capability *per se* but in one or more
20 embodiments, a new piece of code, namely, snapshot management system 2340, is provided. Referring to the exemplary details in FIG. 28, system 2340 has logic to determine when to issue commands. In one or more embodiments, snapshot management system 2340 has a data structure or mapping table or the like that can map a given VM ID into a long-term snapshot. System 2340 also has logic such that when asked to roll back
25 to a long term snapshot, system 2340 will know from the mapping which version of the long term snapshot is to be obtained from the image library, checked out, and swapped as shown in FIG. 25. Snapshot management system 2340 is thus a piece of software that implements the logic in FIGS. 23 and 25 based on external input. As seen in FIG. 23, inputs include a command to take a long-term snapshot, a command to take a short-term
30 snapshot, or a command to roll back to a previous state. API manager 2802 validates the arguments of input commands and creates and connects to any required services.

Orchestrator 2804 decomposes the commands as needed to a sequence of commands that can be understood by the pertinent hypervisor and virtual image library (in this non-limiting example, VMware vCenter and IBM VIL, respectively). Orchestrator 2804 also orchestrates command execution and retrieves and processes results and errors.

5 For the taking of a short-term snapshot, orchestrator 2804 invokes VMware program API 2806 or other VMware adapter after the appropriate checking and conversion to an object are carried out. For the taking of a long-term snapshot, virtual image library query executor 2808 is employed to interact with image library 2202 as shown in FIG. 25. VMDK swap adapter 2810 builds and executes the VMDK swap
10 command to implement the swapping process shown in FIG. 25.

Thus, the step of creating the long-term snapshots is carried out by the VIL 2202 under control of snapshot management system 2340. One suitable virtual image library is the IBM SmartCloud Provisioning Virtual Image Library available from International Business machines Corporation, Armonk, NY, USA.

15 It will be appreciated that a cloud, such as an MIaaS cloud or other cloud, is one non-limiting example of a standardized virtual environment.

In some cases, the creating of the long-term snapshots from time-to-time includes creating the long-term snapshots prior to significant adjustments (e.g., 2222, 2226, and 2230 in FIG. 22) during the migration process.

20 Some cases, as best seen in FIG. 25, further include initiating a rollback. This can be done by:

- locating a given one of the image objects 2576 in the image library 2202, as per step 1;
- 25 • checking the given one of the image objects out of the image library as a template 2572, as per step 2;
- placing the template 2572 into the hypervisor operational repository 2468; and
- instantiating the template from the hypervisor operational repository as a rolled-back virtual machine instance in the hypervisor (where it replaces an existing VM
30 instance 2574 that was problematic, so that rollback to a previous state was desired), as per steps 3 and 4 and locations 2573, 2574.

As noted, in one or more embodiments, such rollbacks can be initiated during a steady state condition after the migration.

As shown in FIG. 26, in some cases, changes made during the migration can be
5 determined by locating a first given one of the image objects in the image library, as per step 2602. The first given one of the image objects corresponds to one of the long-term snapshots associated with an original customer image of the source system (e.g., 2204). A further step in determining changes includes, as per step 2604, locating a second given one of the image objects in the image library. The second given one of the image objects
10 corresponds to one of the long-term snapshots associated with a customer image of the source system after the significant adjustments are complete (e.g., 2210). In step 2606, the differences between the first and second given ones of the image objects are computed (e.g., with engine 2462 – in some cases, a “diff” feature of the aforementioned IBM SmartCloud Provisioning Virtual Image Library can be employed). An even further
15 step includes making a representation of the differences available to a human operator (e.g., via a suitable graphical user interface (GUI) displaying same on a display 24 or the like). The differences that are computed can be, for example, file-level differences and/or product-level differences.

In some cases, a further step can include assigning a version number to each of
20 the long-term snapshots as shown in FIG. 22; i.e., V.1, V.2, ..., V.n. The result is a version chain of images, each corresponding to significant adjustments made; effectively providing a source control system.

In another aspect, an apparatus includes a memory (e.g., RAM 30, cache 32); and at least one processor 16, coupled to the memory, and operative carry out or otherwise
25 facilitate any one, some, or all of the aforementioned method steps. Optionally, the apparatus further includes a plurality of distinct software modules 42. Each of the distinct software modules is embodied on a computer-readable storage medium, and the distinct software modules include a snapshot management system module, a hypervisor module, and a virtual image library module, as discussed elsewhere herein. The snapshot
30 management system module optionally includes one, some, or all of the sub-components in FIG. 28.

REPLACING VIRTUAL MACHINE DISKS

A virtual machine (VM) is a software implementation of a computer that executes programs like a physical machine. Virtual machines may have one or more virtual disks associated therewith. One or more embodiments advantageously facilitate the growth of cloud computing by providing one or more of:

- techniques to migrate one or more disks associated with a source virtual machine (image or instance) into a target virtual machine;
- techniques to restore a virtual machine using a different set of virtual disks; and/or
- techniques that can accomplish migration and/or restoration without the need to alter the unique IDs associated with a virtual machine.

Advantageously, one or more embodiments provide techniques for replacing the disks associated with an existing virtual machine with disks of another virtual machine, or images thereof, without impacting functionality or the identity of the virtual machine. Indeed, one or more embodiments provide techniques to support the ability to replace the disks associated with a virtual machine with the disks associated with another virtual machine or image; techniques which determine the compatibility of the virtual machine configurations and their disks; and/or techniques which ensure that the resulting virtual machine continues to function after the disks have been replaced.

Furthermore in this regard, and referring now to FIG. 34, in at least some embodiments, there are four states overall, as listed below. An “image” is offline / dormant / static, as in “image” of a virtual machine. An instance is a running virtual machine or registered to be able to run in a virtual machine, as in an “instance” of an image. As in FIG. 34, a “virtual machine image” or “image” or “virtual machine template” 3402 includes a virtual disk 3406 containing at least an operating system, and virtual hardware suggestions 3408 such as CPU, memory, and/or disk. A “virtual machine” or “instance” or “instance of an image” 3404 includes 3406 and 3408, as well

as virtual hardware allocations 3410 of CPU, memory, and/or disk, and a hypervisor of virtual machine runtime 3412. The aforementioned four states include:

- 5 1. an image which is existing independent of a hypervisor (i.e. file on a file system)
 – Virtual Machine Image
2. an image which is registered to a hypervisor (i.e. template or clone source) -
 Virtual Machine Image
- 10 3. a virtual machine registered to a hypervisor but is not running (i.e. powered off
 virtual machine) – Virtual Machine
4. a virtual machine registered to a hypervisor and is running (powered on virtual
 machine) – Virtual Machine

15

The replacing of the target could occur from any one of these source states. Therefore, in the most general case, the source can, but need not, be an “image” format as in (1) and (2) or it could be transported directly as per (3) or (4).

20

Furthermore, one or more embodiments provide a systematic way to enable a cloud to adopt an existing server without the need to migrate the application to a net new cloud instance; and/or a rule system to capture which attributes of a source and a target virtual machine should be preserved in order to maintain virtual machine functionality and to be undetectable to the parent hypervisor and/or cloud.

25

In at least some cases, virtual machines can be replaced without making it known to external elements, such as, for example, a cloud management stack. This capability allows for substantial changes to be made to a virtual machine (for example, replacing the contents of all disks) without the hypervisor or dependent management stack being aware of the change. One or more embodiments can be used, for example, to migrate an
30 external virtual machine into a managed cloud, thus enabling the cloud to import existing

virtual machines. Furthermore, one or more embodiments can be used to restore a virtual machine to a previous state.

Attention should now be given to the system diagram of FIG. 29. Note source virtual machine 2901 with one or more disks such as disk 1, numbered 2906, and disk 2, numbered 2908. Disks 2906, 2908, 2910, 2912 are all virtual disks. Virtual disks are discussed further in connection with, e.g., FIG. 36. The system also includes a virtual machine disk replacer such as image replacer 2902, which operates in accordance with merger rules 2914. It is desired to replace Disk 1 and Disk 2, 2910, 2912, on target VM 2903 with disks 2906, 2908. The disks 2910, 2912 also represent the after-migration state; that is, prior to migration, Disk 1 2910 and Disk 2 2912 are in an original state and after migration they have effectively been replaced with disks 2906, 2908 respectively. The target virtual machine 2903 is managed by hypervisor 2904 and the overall target cloud environment is managed by cloud management software 2905.

Referring now also to FIG. 30, cloud hypervisor 3001 (same as hypervisor 2904) is a hypervisor which is managed by the desired target cloud; all instances running within the hypervisor 3001 are assumed to be managed by the cloud. Note that in FIG. 29, the source is depicted external to the cloud hypervisor, while in FIG. 30, both source and target are located on the same hypervisor. These scenarios are different but both are valid and possible. The customer image is imported into the cloud environment and is the source 3002 in the technique shown in FIG. 30; cloud image 3003 is the target. As noted elsewhere, an image of source is not always required. Source 3002 is implemented on virtual machine 3004. Virtual disk descriptor 3006 is a meta-file describing the attributes of the virtual disk(s) 3008, which are the actual block allocations where the disk content is stored.

The cloud adopter 3010 implements one or more techniques described herein in accordance with merge rules 2914. Merge rules 2914 in FIG. 30 and merger rules 2914 in FIG. 29 are the same. Element 3010 is the same as element 2902; element 3004 is the same as element 2901; and element 3005 is the same as element 2903. Disks 3008 include one of 2906, 2908. Disks 3009 include one of disks 2910, 2912. Note that FIG. 36 provides additional details on the relationships of virtual disk descriptors to virtual disk extents to virtual disks. The merge rules 2914 provide specific rules about which

attributes of the source virtual disk descriptor(s), the target virtual disk descriptor(s) and process are included in the final target disk descriptor 3007. The merge rules ensure that the resulting virtual machine 3005 executes and that the cloud management layer and hypervisor 3001 cannot notice a difference. The virtual disks 3009 are the actual block
5 allocations where the disk content is stored in the new virtual machine 3005.

With continued attention to FIGS. 29 and 30, consider also FIG. 31, which depicts a process of replacing a set of disks of a VM 2903, 3005 with those of a source virtual machine 2901, 3004. In step 3101, identify the receiving cloud instance 2903, 3003. In one or more embodiments, the target is a virtual machine (instance) that is in a powered-
10 off state. In step 3102, identify the customer source 2901, 3002. In step 3103, compare the virtual resources of VM 2901, 3004 with those of VM 2903, 3005. In step 3104, determine if the resources are compatible; if not, reject the proposed replacement at 3105 and then the process ends at 3111.

Reference should now be had to FIGS. 35 and 36. Note VM image 3502. In a
15 hypervisor there is typically an overall VM description (file) 3504 setting forth all the virtual hardware characteristics of the virtual machine. In some cases, as shown in FIG. 35, all pertinent disk information regarding virtual disks 3506 is contained in this single master file. In other cases, as seen in FIG. 36, the master file 3604 references other files (descriptors) 3610, 3612 which include the details of one or more virtual disks 3614,
20 3616, 3618.

FIG. 36 is thus an example of a virtual machine image or instance which has one level of indirection between the VM resource descriptor and the actual virtual disk. The Virtual Disk Descriptor contains meta-information about one or more backing disks. This example would present two disks to the virtual machine operating system but is
25 realized by three separate virtual disk files.

Returning now to FIG. 31, if the virtual resources are compatible, as shown at the YES branch of decision block 3104, merge the source virtual resource descriptors 3112 (e.g., virtual disk descriptors 3006) and the target virtual resource descriptors 3113 (e.g., virtual disk descriptors 3007) in step 3106. Furthermore in this regard, elements 3007
30 and 3113 represent the target disk descriptor before the replacement occurs. After the merge and replace, element 3007 becomes an updated version which is a combination of

parameters for the source and the target. Part of step 3108 includes the step to write the new target disk descriptor 3007. See FIG. 32.

In decision block 3107, determine if the merge of the virtual resource descriptors in step 3106 was successful; if not, the process ends at 3111. On the other hand, if the
5 merge of the virtual resource descriptors in step 3106 was indeed successful, proceed to step 3108 and replace the virtual disk as shown at 3303 in FIG. 33. In decision block 3109, determine whether the source VM has more disks; if yes, repeat steps 3106-3109 as needed; if not, update the overall instance descriptor in step 3110. Step 3110 is optional and depends on whether the approach of FIG. 35 or FIG. 36 is used by the VMs that are
10 being operated on. Processing ends at 3111.

Furthermore with regard to FIGS. 35 and 36, implementations can differ in that all virtual machine resource information (including disks) can be stored in a single file as in FIG. 35 or it can be separated into many files as in FIG. 36. The approach is dependent on the hypervisor and has to do with the level of decoupling of components
15 (i.e. for the examples with virtual disk descriptors there is a reference in the overall virtual machine definition linking to those descriptors, then more information in the descriptors about the actual characteristics of the underlying disk). Step 3110 is appropriate where the approach of FIG. 36 is utilized.

FIG. 32 presents exemplary details of the merging step 3106. In step 3203, access
20 the source descriptors 3112. In step 3204, access the target descriptors 3113. In step 3205, examine an attribute. In decision block 3206, determine at least whether the attribute in question is critical to the source, based on the merge rules 2912.

In this regard, the merge rules 2912 are specific to the target hypervisor (non-limiting examples include VMware virtualization software available from VMware, Inc., Palo Alto, CA, USA; KVM software (for Kernel-based Virtual Machine)(available from
25 RedHat, Inc. Raleigh, NC, USA), Xen software (available from Citrix Systems, Inc., Ft. Lauderdale, FL, USA), and so on). In step 3205 and decision block 3206, a number of determinations are made as attributes of both the source and target environments are examined. Steps 3206-708 are discussed further below with regard to FIG. 37. With
30 regard to source attributes, it may be determined, for example, that Attribute A of the source must persist to ensure continued virtual machine functioning in the target cloud.

Examples of such source attributes include the Virtual disk input-output (IO) controller, the block size, the number of blocks, and so on. The same are written to the “new” target descriptor (which is a combination of both source and target parameter/values) in step 3207 – see discussion of FIG. 37 below. With regard to target attributes, it may be
5 determined, for example, that Attribute B of the target must persist to ensure that the cloud successfully adopts the virtual machine (nether hypervisor nor cloud notices a difference). Examples of such target attributes include the virtual disk unique identifier. The same are written to the new merged descriptor in step 3208. Processing then flows to step 3210.

10 In decision block 3210, check to see if there are additional attributes; if so, loop back to step 3205 and repeat the process until all the attributes have been examined. If there are no more attributes, the process ends at step 3211 and flow proceeds to step 3107 in FIG. 31.

Considering now FIG. 37, source virtual machine descriptor 3751 includes
15 Parameters 1-4 with respective values A, B, C, and D. Target virtual machine descriptor 3753 includes Parameters 1-3 with respective values X, Y, and Z. As indicated at 3755, the logic in step 3206 is applied for each parameter in the source and target descriptors, i.e., Parameters 1-4. In step 3206, target parameters dominate unless the source parameter is needed to maintain functionality or compatibility. Thus, by default, all
20 target parameters are retained in updated target virtual machine descriptor 3757, but only the needed source parameters are retained. In the example in FIG. 37, the source parameter values for parameters 1, 3, and 4 were necessary while the target value of “Y” for parameter 2 was retained by default because that source parameter was not needed to maintain identity or functionality. It should be noted that target parameters are assumed
25 to all be compatible as they were the initial values for the existing virtual machine which presumably functioned and was properly known to the hypervisor / management layer.

Thus, one or more embodiments provide a system and method to support the ability to replace the disks associated with a virtual machine with the disks associated with another virtual machine or image; a method which determines the compatibility of
30 the virtual machine configurations and their disks; and/or a method which ensures that the resulting virtual machine continues to function after the disks have been replaced.

Referring now to the system diagram of FIG. 33, one or more embodiments consider an existing physical or virtual server 3301 that does not reside within the managed cloud; rather, server 3301 resides in existing source environment 3350. Server 3301 has existing disks 3351. As seen at 3302, server 3301 is converted to an image 3352. This hypervisor-compatible image is then transported into the target cloud 3354. During cloud adoption process 3303, merger and replacement of the cloud provisioned image with the customer image is carried out. This process can, in some instances, be facilitated by migration engineer 3353. The cloud is unaware of any significant change to a virtual machine (here, VM7 numbered 3365) running under its control, as indicated at 3304.

Target cloud 3354 includes a management layer 3355 (which in at least some instances can be similar to layer 64 described above). Cloud user 3367 interacts with management layer 3355. Management layer 3355 manages hardware and software of the cloud environment 3354 including one or more hypervisors such as 3356, 3357, 3358. The hypervisors in turn manage one or more virtual machines 3366; in the example of FIG. 33, Hypervisor 1, numbered 3356, manages VMs 1-3 numbered 3359-3361; Hypervisor 2, numbered 3357, manages VM 4 numbered 3362; and Hypervisor 3, numbered 3358, manages VMs 5-7 numbered 3363-3365.

Thus, one or more embodiments provide a method which allows a non-cloud provisioned image to be adopted by the cloud and be treated as any other cloud provisioned instance; a process which defines how to move an existing physical or virtual server into a cloud without formal support; a system that supports merging virtual machine descriptors; and/or replacing target image disks with source image disks; and/or a method of merging virtual resource descriptors. This latter method can include, for example, retrieving virtual resource descriptor(s) from a source and target image; comparing the attributes of each virtual resource descriptor; and/or merging both descriptors into a third descriptor which is compatible with the target virtual machine configuration, the target hypervisor, the target virtual machine operating system and the target cloud.

One or more embodiments thus provide the capability to replace one or more virtual disks associated with a source virtual machine over the same number of virtual

disks associated with a target virtual machine. This capability enables the ability to migrate either a virtual machine or virtual machine image into an existing virtual machine definition. This is relevant, for example, in a cloud or managed cloud scenario where the identity and features of the virtual machine are managed at a higher level (i.e. to track capacity, licensing, and so on). If the identity changes, the management layer would lose track of the system. Another scenario allows for the restoration of a virtual machine to a previous state by allowing its current set of virtual disks to be replaced by a previously captured set of virtual disks. A further possibility allows the re-use of virtual machine “containers” by simply replacing the virtual disk contents as opposed to creating entirely new containers each time.

Given the discussion thus far, it will be appreciated that, in general terms, an exemplary method, according to an aspect of the invention, includes the step 3106 of merging at least one target virtual disk descriptor 3007, 3113 that describes at least one virtual disk 2910, 2912, 3009 associated with an existing target virtual machine 2903, 3005 in a target virtualized environment with at least one source virtual disk descriptor 3006, 3112 that describes at least one virtual disk 3352 associated with a source. The merging is carried out in order to obtain at least one merged virtual disk descriptor (also represented by 3007; i.e., element 3007 represents the descriptor before the merge and also after the merge at which point it has some source parameters). The at least one merged virtual disk descriptor is compatible with the target virtualized environment. In some cases, step 3106 can be carried out using a “MergeDiskMetaData” sub-module 3913 of a swap module 3901, executing on at least one hardware processor, as described below in connection with FIG. 39. A further step 3108 includes replacing the at least one virtual disk associated with the existing target virtual machine in the target virtualized environment with the at least one virtual disk associated with the source, in accordance with the merged plurality of virtual resource descriptors. In some cases, step 3108 can be carried out using a “PutDisk” sub-module 3915 of a swap module 3901, executing on at least one hardware processor, as described below in connection with FIG. 39.

Furthermore in this regard, in FIG. 39, note swap module 3901. Preparation sub-module 3902 ensures that the source and target VMs are in a state wherein they can be operated upon (image versus instance). Any needed preparations are carried out by

prepSource sub-module 3903 and prepTarget sub-module 3905, for the source and target, respectively. CheckCompatibility sub-module 3907 makes sure that the source and target are compatible. SwapDisks sub-module 3909 and replaceDisk sub-module 3911 include control logic to implement the looping in FIG. 31 through the tasks assigned to sub-modules 3913 and 3915; i.e., steps 3106 and 3108.

TransformDisk sub-module 3917 addresses issues if the source does not match the target specification; it transforms the actual virtual disk to make it indistinguishable. One example is if the source and target virtual disks are of different sizes. This will typically not impact the hypervisors but could cause issues with future operations such as “extend disk.” The updateVMSpec sub-module 3919 carries out the optional step 3110 associated with FIG. 36-type configurations. Hypervisor adapter module 3930 merges in each of the functions to the APIs 3931 exposed by the hypervisors 3932, 3933, 3934; namely, the ability to copy files, move files, obtain descriptors, and update descriptors. The ellipses with regard to 3933, 3934 refer to “operation implementation” as spelled out in 3932.

Furthermore with regard to prepare block 3902, refer to FIG. 38. The prepare function obtains or accesses the source image 3802, source physical machine 3804, or source virtual machine 3806 and brings same to a state where the replace operation can be carried out. Source image 3802 can simply be deployed to a VM under a hypervisor, as shown at 3812. Source physical machine 3804 can be converted to an image in the case of no hypervisor, as shown at 3808, or converted to a VM in the case of a hypervisor, as shown at 3810. Source VM 3806 can also be converted to a VM as in 3810. In each case, flow proceeds as appropriate to the target VM disk replacement at 3814.

The modules and/or sub-modules shown in FIGS. 38 and 39 can be written in a variety of languages; Perl is a non-limiting example.

In some cases, the target virtualized environment is a target cloud environment. One or more embodiments are believed to be especially useful with respect to clouds because clouds are sensitive to the changing of identity.

One or more embodiments thus address the replacement of virtual disks which is supported by the merging of meta-data about those virtual disks.

In some instances, the source is a customer source outside the target cloud environment, and the replacing includes migrating the at least one disk associated with the source to the at least one disk associated with the existing target virtual machine, as depicted, e.g., in FIG. 33.

5 In at least some embodiments, an additional step 3103 includes comparing virtual resources of the at least one disk associated with the existing target virtual machine and the at least one disk associated with the customer source outside the target cloud environment. In such cases, the merging is in response to the comparing of the virtual resources indicating compatibility thereof, as per the YES branch of decision block 3104.

10 In at least some embodiments, an additional step 3107 includes checking for success of the merging step. In such cases, the replacing is in response to the checking indicating the success of the merging step, as per the YES branch of decision block 3107.

With regard to the decision in decision block 3104, one or more embodiments make use of a data repository 3199 which tracks known compatibilities. These compatibilities, for instance, would compare the IO adapter or the like across the source and target to determine if the source operating system would be able to function with the target adapter (if different). In some embodiments, a check is also made to ensure that the number of disks matches (i.e. might be unable to merge a source with two disks into a target with one disk). Overall, the content of database 3199 may, for example, be mined from the technical limitations associated with both the operating system and software contained within the virtual machine / image as well as any compatibility issues at the hypervisor level (i.e. maximum disk size, maximum number of disks, and so on).

20 An analogous approach may be used to develop a data store (omitted from figure to avoid clutter) on which the decision in block 3107 may be based.

25 As indicated by decision block 3109, in some cases, the merging and replacing steps are repeated for one or more additional source virtual disks. This overall flow can operate on each virtual disk associated with the source image / virtual machine. The decision block 3104 check is a go/no-go before any changes are made. The 3106, 3108 processes are the actual updating of the disk meta-data and replacing of the disk.

In at least some cases, an additional step 3110 includes updating an overall instance descriptor. Again, step 3110 is optional and depends on whether the approach of FIG. 35 or FIG. 36 is used by the VMs that are being operated on.

In some cases, the merging step 3106 includes, as seen at 3205, examining at least one source attribute associated with the at least one source virtual disk descriptor and examining at least one target attribute associated with the at least one target virtual disk descriptor. A further step 3206 includes applying a plurality of merge rules 2912 to the at least one source attribute and the at least one target attribute to obtain the at least one merged virtual disk descriptor. The merge rules persist, in the at least one merged virtual disk descriptor, source attributes needed to ensure continued virtual machine functioning in the target cloud and target attributes required for the target cloud to successfully adopt the virtual machine. Furthermore in this regard, in one or more embodiments, there are descriptors for the source, the target before the merge, and the target after the merge. The source and target merge to create an updated target. The target virtual machine continues to function properly as does the runtime of the virtual disks which have been replaced into it.

Non limiting examples of merge rules used in applying step 3206 include persisting a virtual disk input-output controller of the at least one virtual disk associated with the source; persisting a block size and number of blocks of the at least one virtual disk associated with the source; and persisting a virtual disk unique identifier of the at least one virtual disk associated with the existing target virtual machine.

In some cases, a source disk image is already available. Also, as discussed above, the replacing of the target could occur from any one of source states (1)-(4), and in the most general case, the source can, but need not, be an “image” format. In some cases, as seen at 3302, an additional step includes converting a source instance to obtain a source disk image 3352.

In at least some cases, the target virtual machine 2903, 3005 has a target virtual machine configuration; the target virtual machine 2903, 3005 is hosted on (a guest of) a target hypervisor 2904, 3001; and the target virtual machine 2903, 3005 optionally has a target operating system. In the merging step 3106, the compatibility with the target virtualization environment includes at least compatibility with the target virtual machine

configuration and the target hypervisor. Optionally the compatibility with the target virtualization environment further includes compatibility with the target operating system.

5 With regard to the target virtual machine not having an operating system in some cases, one example is a PXE boot system or system booting from CD/DVD; the same could have all of its disks replaced without impacting the “operating system.”

10 In addition to migration applications, in some cases, the source 2901 is a back-up source; and the replacing includes restoring the at least one disk (e.g., 2910 or 2912) associated with the existing target virtual machine using the at least one disk image associated with the back-up source.

15 As noted elsewhere, in one or more embodiments, there is no change to functionality and identity of the VM 2903, 3005, 3365, and/or the process is not detectable to the parent hypervisor 3358 and cloud management 3355. It should be noted that hypervisors typically have the concept of a unique identifier for the virtual machine that is not duplicated within the virtualization environment. The skilled artisan will be familiar with same, for example, from VMware and KVM (for Kernel-based Virtual Machine, a full virtualization solution for Linux on x86 hardware containing virtualization extensions). The VMware Virtual Machine UUID (backed by a field called uuid.bios in the virtual machine descriptor file – a .vmx) is a non-limiting example of a
20 unique identifier. This, in one or more embodiments, any type of unique identifier or associated data stored in the hypervisor or management stack remains consistent and unchanged. The operation / behavior of the replaced target disks is equivalent to that of the original runtime with the source disks. Thus, in some cases, when replacing the at least one virtual disk associated with the existing target virtual machine in the target
25 virtualized environment with the at least one virtual disk associated with the source, a unique identifier of the existing target virtual machine is maintained prior and subsequent to the replacement.

30 In another aspect, an apparatus includes a memory (e.g., RAM 30, cache 32); and at least one processor 16, coupled to the memory, and operative carry out or otherwise facilitate any one, some, or all of the aforementioned method steps. Optionally, the apparatus further includes a plurality of distinct software modules 42. Each of the

distinct software modules is embodied on a computer-readable storage medium, and the distinct software modules include a MergeDiskMetaData module and a PutDisk module; as discussed elsewhere herein.

5 ADJUSTMENT TO MANAGED-INFRASTRUCTURE-AS-A-SERVICE CLOUD
STANDARD

A hardware infrastructure as a service cloud (HIaaS Cloud) provides bare bone virtual machines as a service. It may also provide an operating system and applications; however, support is typically not provided for any software installed on the virtual
10 machine (VM). A non-limiting example is the Amazon Elastic Compute Cloud (Amazon EC2) available from Amazon Web Services LLC Seattle, Washington, USA. A managed infrastructure as a service cloud (MIaaS Cloud) provides full service virtual machines. The service contract includes providing software services for the underlying virtual machine. These services include software upgrades, patch support, support for
15 compliance, license management, and the like for the provided software infrastructure. A non-limiting example is the IBM SmartCloud Enterprise+ also known as IBM SCE+, cloud environment available from International Business Machines Corporation, Armonk, NY, USA. Image-based migration uses customer images for migration of the application to the cloud. This approach minimizes expensive application re-installation.
20 A cloud standard is the minimum set of requirements needed to be met by a customer image to be run and managed in cloud.

One or more embodiments advantageously provide a system and method for rapid, reliable, and reduced cost adjustment to MIaaS cloud standard. One or more embodiments take an arbitrary customer instance and perform a set of changes so that the
25 MIaaS management layer can manage the customer application exactly as it would manage an application created from its own golden masters (reference models from which copies are mass-produced). One or more embodiments provide a process which is rapid (the end-to-end duration of standardization is small (e.g., < 4 hours)); exhibits reduced cost (the amount of time spent by a migration engineer is quite small (e.g., < 30
30 minutes); and reliable (images coming out of the process are standard with high reliability and meet all customer preferences). Furthermore, one or more embodiments

provide a two-step process wherein customers and cloud engineers formulate a flexible standardization plan in line with customer preferences and cloud standard, and wherein an automated framework executes the plan in a reliable fashion in conformance with the rules.

5 In one or more embodiments, rules are used to uniformly encode adjustment recipes across various kinds of images; image adjustments are used for standardization; and/or a generic orchestration framework is used to deal with multiple types of adjustment actions.

10 Currently, with regard to attempts to standardize migration to MIaaS environments, it should be noted that an MIaaS delivery model is dependent on having pre-fabricated virtual machine images, which meet all the services requirements. Currently, standardization for image-based migration is manual and suffers from being time consuming, costly, and unreliable. Current techniques are time-consuming because a large number of long running steps are needed. Current techniques are costly because
15 all steps are performed manually. Furthermore, current techniques are unreliable; furthermore in this regard, since source instances are different, a migration engineer needs to manually identify the right set of steps to be performed on each source instance. The result is that required steps may be missed; unwanted steps may be performed and corrupt the instance; and/or many decisions are not technical and need to be taken in
20 consultation with the customer. That is to say, the migration engineer may not be suited for deciding on the fly.

One or more embodiments advantageously provide a rapid standardization approach. One or more embodiments define a generic framework, as seen in FIG. 40, in which standardization plans can be injected. Note adjustment framework 4002;
25 supporting tools 4004 including individual tools 4020, 4022, 4024; and adjustment zone 4006. Adjustment plans 4018 can encompass a wide variety of adjustment actions. Each type of action is orchestrated by its orchestrator 4012, 4014, 4016; the orchestrators operate under control of adjustment choreographer 4008. New actions can be added to an existing plan, allowing the framework to be extensible. Adjustment actions can also
30 include offline adjustments, which do not require an instance to be running.

New orchestrators 4012, 4014, 4016 and actions can be registered in the adjustment plan 4018 (for example, using a suitable plan manager 4010). Each orchestrator 4012, 4014, 4016 may interact with the virtual machines 4030, 4032, 4034 (online), image repository 4028 (offline), virtualization manager 4026, or any other supporting tools (e.g., IBM Tivoli® Endpoint Manager (TEM) server, IBM Tivoli Provisioning Manager available from International Business Machines Corporation, Armonk, NY, USA; VMware vCloud Director available from VMware, Inc. Palo Alto, CA USA).

One or more embodiments provide a rapid standardization system and method which utilizes offline adjustments. Offline adjustments are file system patches that can be applied to a virtual machine image. Adjustments will make changes to the image to meet the cloud standard (e.g., install a Secure Shell (SSH) service). Offline adjustments can be applied in parallel and without using up any compute resources, thus enabling rapid adjustments. Offline adjustments include simple file system copies or other operations not requiring the VM to be activated. This takes much less time than actual execution of the install/configuration steps (again, enabling rapid adjustment). Offline adjustments can be applied without booting up the image and do not need network access to the image. This advantageously solves the chicken-and-egg problem wherein access to the VM in cloud requires configuration (thus removing manual steps).

The flow chart of FIG. 41 shows the flow of a sample offline adjustment. The flow chart begins at 4102. In decision block 4104, determine, for example, whether cygwin/ssh service exists. In this regard, Cygwin is a Unix-like environment and command-line interface for Microsoft Windows. Cygwin provides native integration of Windows-based applications, data, and other system resources with applications, software tools, and data of the Unix-like environment. Thus it is possible to launch Windows applications from the Cygwin environment, as well as to use Cygwin tools and applications within the Windows operating context.

In this example, as seen at the “YES” branch of block 4104, if the cygwin/ssh service already exists, further processing is not needed so proceed to the end and 4106. On the other hand, if the cygwin/ssh service does not already exist, as seen at the “NO” branch of block 4104, in step 4108, power of the virtual machine and mount the virtual

machine. In step 4110, copy the cygwin manifest onto the image. In step 4112, copy the required public keys onto the image. In step 4114, update the registry to run the script on the first boot of an instance. In step 4116, dismount the virtual machine and power on the virtual machine. In step 4118, the script to enable the ssh service now executes automatically; proceed to end step 4106. It should be emphasized that FIG. 41 is but one specific detailed example demonstrating how an offline adjustment can be carried out in one or more embodiments. Furthermore in this regard, as noted, offline adjustments are file system patches that can be applied to a virtual machine image; implementing cygwin/ssh service is thus a non-limiting specific example of many kinds of offline adjustments that can be carried out in one or more embodiments.

Reference should now be had to FIGS. 42 and 43. One or more embodiments provide reduced cost standardization via an automated adjustment engine 4302. One or more embodiments include an adjustment orchestrator 4304 and/or 4310. The orchestrator 4304 and/or 4310 defines a set of phases where a dependent phase is executed only after all its antecedent phases have been executed. Steps in each phase are encapsulated as a script. Scripts can be run on the adjustment engine 4302, on target virtual machines, or on any other management tool (e.g., integrated with vCenter (VMware vCenter Server provides a centralized and extensible platform for managing virtual infrastructure, and is available from VMware, Inc. of Palo Alto, CA, USA) or TEM). In one or more embodiments, a single user interface 4302 is provided for all kinds of images (e.g., Windows, Linux, AIX). Furthermore in this regard, in an exemplary embodiment, element 4302 is the adjustment portal and the master choreographer 4008, which decides the order in which the individual orchestrators 4012, 4014, 4016 will be invoked. The adjustment engine is the sum of all components. The portal is an optional entity in element 4302 as the choreographer can be invoked directly using a programmable API. In summary, in the non-limiting example, element 4302 equals the choreographer plus, optionally, the portal. In one or more embodiments, multiple images can be adjusted in parallel (this reduces the amount of time spent per image being standardized).

One or more embodiments provide reliable standardization. One or more embodiments provide a flexible rule framework that takes conditions and actions.

Conditions provide the flexibility to handle different types of source images. Conditions and actions come up with associated scripts. In at least some cases, this involves a two-step process. The customer approves the valid set of rules for his or her data center and bootstraps the Adjustment Orchestrator 4304 and/or 4310 (see also 4012-4016) with the rules. The tool allows migration engineers to execute the standardization in conformance with the rules. A log is maintained to audit the standardization process any later time. Furthermore, in one or more embodiments, snapshots are employed to allow rollback to a reliable version of the image in case of standardization failures. In some embodiments, at least a two-step process is employed; however, in certain cases a multiple step process is utilized. The customer may define a set of migration rules and the migration engineer may provide feedback that some rules cannot be enforced and thus, there could be an iterative back-and-forth.

The following is an exemplary, illustrative rule set that may be used in the reliable standardization process described just above:

- If (Cygwin not installed) Then (Install Cygwin)
- If (ITM (IBM Tivoli Monitoring software, available from International Business Machines Corporation, Armonk, NY, USA) is present) Then (Uninstall ITM)
- If (True) Then (Insert ssh keys)
- If (Symantec Antivirus is not present) Then (Install Symantec Antivirus)
- If (Password Expiration = 90 days) Then (Set Password Expiration = 180 days)
- If (Patch X not present) Then (Deploy Patch X) (Patch X here represents any needed patch; a non-limiting example could be, for example, Microsoft Patch MS08-067)
- If (TEM agent not registered) Then (Register TEM Agent)

Refer now to FIG. 42. At 4202, mount the imported image. In step 4224, carry out access enablement; for example:

1. Add Cygwin SSHD (OpenSSH = sshd) for Windows;
2. Add SSH keys;

3. Set IP addresses and default gateway to all network interfaces on first boot;
4. Set name server entries; and
5. Take a snapshot.

5 In the event of any failure, as seen at 4204, fix the errors, rollback, and retry. Once all steps are successful, whether on the first attempt or a subsequent re-try, as seen at 4206, start the VM. Then implement an agent un-install procedure 4226; for example:

1. Uninstall TEM agent
- 10 2. Uninstall EMC Networker (available from EMC Corporation, Hopkinton, MA, USA)
3. Uninstall TSM client (IBM Tivoli Storage Manager (TSM) software available from International Business Machines Corporation, Armonk, NY, USA)
4. Uninstall TSCM agent (IBM Tivoli Security Compliance Manager (TSCM)
- 15 software available from International Business Machines Corporation, Armonk, NY, USA)
5. Uninstall TAD4D agent (Tivoli Asset Discovery for Distributed, available from International Business Machines Corporation, Armonk, new York, USA)
6. Uninstall ITM agent (IBM Tivoli Monitoring software, available form
- 20 International Business Machines Corporation, Armonk, NY, USA)
7. Take snapshot

 In the event of any failure, as seen at 4208, fix the errors, rollback, and retry. Once all steps are successful, whether on the first attempt or a subsequent re-try, as seen

25 at 4210, proceed to the agent install process 4228; for example:

1. Install TEM
2. Set Domain Name System (DNS) Cache Manual

In the event of any failure, as seen at 4212, fix the errors, rollback, and retry. Once all steps are successful, whether on the first attempt or a subsequent re-try, as seen at 4214, proceed to the patching process 4230; for example:

- 5 1. Patch to baseline
2. Take snapshot

In the event of any failure, as seen at 4216, fix the errors, rollback, and retry. Once all steps are successful, whether on the first attempt or a subsequent re-try, as seen
10 at 4218, proceed to the compliance process 4232; for example:

1. Set Minimum Password Length;
2. Set Password History Count
3. Set Maximum Password Age
- 15 4. Set Minimum Password Age
5. Set Password Complexity
6. Take a snapshot.

In the event of any failure, as seen at 4220, fix the errors, rollback, and retry.
20 Once all steps are successful, whether on the first attempt or a subsequent re-try, as seen at 4222, proceed to the on-boarding process.

FIG. 43 shows exemplary standardization architecture. Included are adjustment engine 4302 (which may be provided as a graphical user interface or GUI or a set of APIs to a master choreographer); and online and offline adjustment orchestrators 4304, 4310, respectively (both of which are in the form of components). Other components include
25 adjustment database 4306, adjustment plan maintenance module 4308, access manager 4312, compliance manager 4316, backup manager 4318, script execution manager 4320, and patch manager 4322. Engine 4302 invokes the online and offline adjustment orchestrators 4304, 4310. These in turn query adjustment database 4306, which is
30 maintained by using adjustment plan maintenance module 4308. The adjustment plan maintenance component is used by the migration team in consultation with the customer

to create an adjustment plan in advance. This plan contains all the rules to guide the adjustment process.

Offline adjustment orchestrator 4310 uses access manager 4312 to provide required access and interfaces.

5 Online adjustment orchestrator 4304 invokes compliance manager 4316, backup manager 4318, and script execution manager 4320. Compliance manager 4316 also invokes script execution manager 4320. Online adjustment orchestrator 4304 uses patch manager 4322.

10 Attention should now be given to FIGS. 44-46, which show exemplary detailed steps in an adjustment phase. Each step is shown in a “swim lane” indicating what module or entity carries out the particular step in the exemplary embodiment. Please note that elements 4402, 4408 (cloud portal), and 4410 are not needed for the process depicted in FIG. 44 but are included for illustrative purposes. Migration team 4404 captures all activities performed by the migration team. This includes activities that are performed by
15 using the Adjustment GUI or invoking the Adjustment APIs as well as manual steps to check for success and failure. Test and Troubleshoot VM steps 4432, 4442 represent a manual process in the example of FIG. 44. The remaining steps are typically automated and involve clicking a button on the portal or calling an API. The entities include the TSAM (IBM Tivoli® Service Automation Manager software available from International
20 Business Machines Corporation, Armonk, NY, USA) administrator 4402, migration team 4404, VMware vSphere Virtualization Manager 4406, Portal 4408, TSAM 4410, and back-up library 4412. Processing begins at 4420. In step 4422, carry out smoke testing (i.e., a process of validating code changes before the changes are checked into the product’s source tree). In step 4424, back up the VM in question; in step 4426, check
25 same into the back-up library. In step 4428, the hypervisor provides the VM; in step 4430, configure the VM. Step 4430 is analogous to step 4224 in FIG. 42. In step 4432, test and troubleshoot the VM. In step 4434, back up the configured and tested VM; in step 4436, check same into the back-up library. In step 4428, the hypervisor provides the configured and tested VM; in step 4440 remove inappropriate software and/or agents.
30 Step 4440 is analogous to step 4226 in FIG. 42. In step 4442, test and troubleshoot the

updated VM with the inappropriate software and/or agents removed. Processing continues at 4444.

In FIG. 45, processing begins at point 4446 (which is identical to point 4444). In step 4448, back up the updated VM with the inappropriate software and/or agents removed; in step 4450, check same into the back-up library. In step 4452, the hypervisor provides the updated VM with the inappropriate software and/or agents removed; in step 4454 install appropriate agents. Step 4454 is analogous to step 4228 in FIG. 42. In step 4456, test and troubleshoot the updated VM with the appropriate agents installed. In step 4458, back up the updated VM with the appropriate agents installed; in step 4460, check same into the back-up library. In step 4462, the hypervisor provides the updated VM with the appropriate agents installed; in step 4464 install patches. Step 4464 is analogous to step 4230 in FIG. 42.

In step 4466, test and troubleshoot the updated VM with the patches installed. In decision block 4468, determine if the testing was successful; if yes, at 4470, proceed to the flow in FIG. 46; if no, at 4472, fix errors, rollback, and retry, as at 4216 in FIG. 42. In FIG. 46, processing begins at point 4474 (which is identical to point 4470). In step 4476, back up the updated VM with the patches installed; in step 4478, check same into the back-up library. In step 4480, the hypervisor provides the updated VM with the patches installed; in step 4482 ensure compliance. Step 4482 is analogous to step 4232 in FIG. 42. In step 4484, test and troubleshoot the post-compliance-check VM. In decision block 4486, determine if the testing was successful; if yes, at 4488, proceed with on-boarding; if no, at 4490, fix errors, rollback, and retry, as at 4220 in FIG. 42. In general, in case of errors, roll back to any backup step in the flow is possible.

FIG. 47 shows an exemplary architecture, to assist the skilled artisan in understanding an exemplary system in a pertinent context. One or more adjustments 4702 are carried out to standardize an arbitrary customer instance. Prior to attempting adjustments, it is preferable to place a snapshot in a suitable image library such as IBM Tivoli Image Library 4718. This can be carried out, for example, using hypertext transfer protocol (http) on port 9443 using IBM Tivoli Image Library. Reference can be had to scope database 4704 to access pertinent data regarding VMs and applications, and to determine the status of VMs. This can be carried out, for example, using Java Database

Connectivity (JDBC). The image to be adjusted is mounted in the network file system (NFS) storage 4706. Any required VM actions are carried out by cloud vCenter 4710, controlled, by vCenter driver 4712; this results in creation and/or modification of VMs in cloud 4708 which are in turn mounted in NFS storage 4706.

5 An IBM Tivoli® Endpoint Manager (TEM) server 4716 may be employed for migration purposes; communication with same may be had, for example, via transfer control protocol (tcp). The TEM controls VM patching for the cloud managed environment 4714 in which adjusted VMs reside; communication can be via http or User Datagram Protocol (udp). TEM 4716 communicates with Internet 4720 via http on port
10 80.

One or more embodiments thus provide a system and method to adjust images so that they meet the image standard of a managed environment 4714 and can be managed using standard processes and tools. In one or more embodiments, such a system includes one or more adjustment type orchestrators 4012, 4014, 4016 handling specific types of
15 adjustments; an adjustment plan 4018 capturing all the adjustments; and an adjustment plan maintenance component 4010.

Non-limiting examples of orchestrators include an Access Enablement Orchestrator 4224 to ensure that all source images can be accessed by all relevant management entities (including other orchestrators); an Uninstall Orchestrator 4226 to
20 uninstall unwanted software; an install orchestrator 4228 to install required software; a Patch Orchestrator 4230 to ensure patch standardization; and a compliance orchestrator 4232 to ensure compliance with one or more cloud standards. Different embodiments could include additional orchestrators, fewer orchestrators, or precisely the orchestrators just described, for example.

25 The managed environment 4714 may be, for example, an MIaaS cloud (e.g., IBM SCE+).

The orchestrators may operate on file system virtual machine images (e.g., in image repository 4028) or on VM instances 4030, 4032, 4034.

In some cases, the actions executed by each of the orchestrators are determined by
30 a rule engine. Furthermore in this regard, the rule engine is not shown explicitly in the figures to avoid clutter; however, in one or more embodiments a rule evaluation engine is

implicitly implemented in each orchestrator, which needs to evaluate each condition and execute the action only if the condition is satisfied. In some cases, such a rule engine includes a set of condition-action pairs (e.g., If (Cygwin not installed) Then (Install Cygwin) and so on, as discussed above); in such cases, the action is executed only when
5 the condition is satisfied. In some cases, new rules can be injected in the rule engine to extend the tool for dealing with new kinds of customer images; for example, such rules can be injected by an adjustment plan maintenance component such as plan manager 4010, which allows for creation, editing, and/or deletion of rules.

10 In at least some cases, the overall adjustment flow can be tuned by changing the adjustment plan 4018 and/or changing the order of the orchestrators 4012-416 or 4224-632.

Recapitulation

Given the discussion thus far, it will be appreciated that, in general terms, an
15 exemplary method, according to an aspect of the invention, includes the step transferring an arbitrary customer instance (generally, one or more such instances) from a customer environment to a target infrastructure-as-a-service cloud environment (4714; FIG. 3) as a (generally, one or more) transferred image. “Transferring” should be broadly understood to include acts by the provider of the infrastructure-as-a-service cloud environment
20 and/or acts by the customer. This step can be carried out, for example, by a transfer core component (e.g., a module 42 in the target infrastructure-as-a-service cloud environment configured to obtain the arbitrary customer instance from a customer environment). A further step includes formulating an image adjustment plan 4018 capturing at least one adjustment required to render the transferred image(s) compliant with standards of the
25 target infrastructure-as-a-service cloud environment. “Formulating” should be broadly understood to include acts performed solely by a machine (e.g., storing instructions, input by a human, which codify the plan) or acts by both a machine and a human. This step can be carried out, for example, by adjustment plan maintenance component 4010. A further step includes executing the image adjustment plan to adjust the transferred
30 image(s) to obtain an adjusted image compliant with standards of the target infrastructure-as-a-service cloud environment. This step can be carried out by an

orchestration component; e.g., choreographer 4008 and one or more orchestrators 4012-4016. A still further step includes on-boarding the adjusted image(s) into the infrastructure-as-a-service cloud environment as a standard image thereof (generally, one or more standard image(s)). “Standard” implies that the image then works in the same
5 manner as if it had been made from the systems own “golden masters.” This step can be carried out, for example, by cloud management layer 64.

In some cases, the transferred image is adjusted online while running as an instance in the infrastructure-as-a-service cloud environment (e.g., 4030, 4032, 4034). In such cases, the executing step includes executing the image adjustment plan with an
10 automated framework including at least one online orchestrator 4304 interacting with the instance in the infrastructure-as-a-service cloud environment.

In some cases, the transferred image is adjusted offline in a file system 4028 of the infrastructure-as-a-service cloud environment. In such cases, the executing step includes executing the image adjustment plan with an automated framework including at
15 least one offline orchestrator interacting 4310 with the transferred image in the file system of the infrastructure-as-a-service cloud environment.

In some cases, the executing step includes executing the image adjustment plan with an automated framework including at least one orchestrator interacting with a virtualization manager 4026 (and/or supporting tools 4004) of the infrastructure-as-a-
20 service cloud environment.

In at least some embodiments, the infrastructure-as-a-service cloud environment includes a managed infrastructure-as-a-service cloud environment having a management layer 64. In such cases, in at least some instances, a further step includes managing the standard image with the management layer together with at least one application created
25 from a managed infrastructure-as-a-service cloud environment application golden master (i.e., an application created from a golden master associated with the managed infrastructure-as-a-service cloud environment; not meant to imply that it is somehow a golden master of the whole cloud environment).

In some embodiments, the image adjustment plan is executed with an automated
30 framework including at least one of an access enablement orchestrator 4224 which ensures that the adjusted image can be accessed by a management layer of the target

infrastructure-as-a-service cloud environment; an agent uninstall orchestrator 4226 which uninstalls unwanted software from the transferred image; an agent install orchestrator 4228 which installs required software to the transferred image; a patching orchestrator 4230 which ensures that the adjusted image has installed thereon patches required by the standards of the target infrastructure-as-a-service cloud environment; and a compliance orchestrator 4232 which ensures that the adjusted image conforms to the standards of the target infrastructure-as-a-service cloud environment.

In some cases, the executing step includes executing the image adjustment plan with an automated framework by using an access enablement orchestrator 4224 to ensure that the adjusted image can be accessed by a management layer of the target infrastructure-as-a-service cloud environment; using an agent uninstall orchestrator 4226 to uninstall unwanted software from the transferred image; using an agent install orchestrator 4228 to install required software to the transferred image; using a patching orchestrator 4230 to ensure that the adjusted image has installed thereon patches required by the standards of the target infrastructure-as-a-service cloud environment; and using a compliance orchestrator 4232 to ensure that the adjusted image conforms to the standards of the target infrastructure-as-a-service cloud environment.

In some cases, adjustment plan 4018 includes rules that uniformly encode adjustment recipes across different kinds of images. Preferably but optionally, the orchestration framework is generic and can deal with multiple types of adjustment actions. In at least some cases, the adjustment orchestrator(s) define(s) a set of phases with steps in each phase encapsulated in a script. Such scripts may run in an adjustment engine and may target VMs 4030-4034, management tools 4020-4024, or the like. Some techniques adjust multiple images in parallel. In some cases, the adjustment orchestrator(s) is/are bootstrapped with rules. The actions of the orchestrators are determined by a rule engine (a set of condition/action pairs); new rules can be injected by adjustment plan maintenance component 4010. The adjustment flow can be changed by changing the plan and/or changing the order of the orchestrators. The framework can be extended by adding and/or registering a new action and/or orchestrator to an existing plan.

In another aspect, an exemplary apparatus includes a memory (e.g., RAM 28, cache 32); at least one processor 16, coupled to the memory; and a non-transitory computer-readable storage medium 34, embodying in a non-transitory manner a plurality of distinct software modules 42. The plurality of distinct software modules in turn include an adjustment plan maintenance component module 4010 including instructions which, when loaded into the memory, configure the at least one processor to formulate an image adjustment plan 4018 capturing at least one adjustment required to render a transferred image compliant with standards of a target infrastructure-as-a-service cloud environment. The transferred image includes an arbitrary customer instance from a customer environment. Also included are an adjustment orchestration module (e.g., 4008 and 4012-4016) including instructions which, when loaded into the memory, configure the at least one processor to execute the image adjustment plan to adjust the transferred image to obtain an adjusted image compliant with standards of the target infrastructure-as-a-service cloud environment; and a cloud management module 64 including instructions which, when loaded into the memory, configure the at least one processor to on-board the adjusted image into the infrastructure-as-a-service cloud environment as a standard image thereof.

In some cases, the adjustment orchestration module includes instructions which, when loaded into the memory, configure the at least one processor to adjust the transferred image online while running as an instance in the infrastructure-as-a-service cloud environment, as at 4030-4034. In some cases, the adjustment orchestration module includes instructions which, when loaded into the memory, configure the at least one processor to adjust the transferred image offline in a file system 4028 of the infrastructure-as-a-service cloud environment.

In some embodiments, the adjustment orchestration module includes instructions which, when loaded into the memory, configure the at least one processor to interact with a virtualization manager 4026 of the infrastructure-as-a-service cloud environment. In some embodiments, the adjustment orchestration module includes instructions which, when loaded into the memory, configure the at least one processor to interact with a supporting tool 4020-4024 of the infrastructure-as-a-service cloud environment.

In at least some cases, the infrastructure-as-a-service cloud environment includes a managed infrastructure-as-a-service cloud environment 4714 having a management layer 64. The management layer can, for example, manage the standard image together with at least one application created from a managed infrastructure-as-a-service cloud environment application golden master.

In some embodiments, the adjustment orchestration module includes an adjustment choreographer 4008 and at least one of: an access enablement orchestrator 4224 which ensures that the adjusted image can be accessed by a management layer of the target infrastructure-as-a-service cloud environment; an agent uninstall orchestrator 4226 which uninstalls unwanted software from the transferred image; an agent install orchestrator 4228 which installs required software to the transferred image; a patching orchestrator 4230 which ensures that the adjusted image has installed thereon patches required by the standards of the target infrastructure-as-a-service cloud environment; and a compliance orchestrator 4232 which ensures that the adjusted image conforms to the standards of the target infrastructure-as-a-service cloud environment.

In some cases, the adjustment orchestration module includes an adjustment choreographer 4008; an access enablement orchestrator 4224 which ensures that the adjusted image can be accessed by a management layer of the target infrastructure-as-a-service cloud environment; an agent uninstall orchestrator 4226 which uninstalls unwanted software from the transferred image; an agent install orchestrator 4228 which installs required software to the transferred image; a patching orchestrator 4230 which ensures that the adjusted image has installed thereon patches required by the standards of the target infrastructure-as-a-service cloud environment; and a compliance orchestrator 4232 which ensures that the adjusted image conforms to the standards of the target infrastructure-as-a-service cloud environment.

Exemplary System and Article of Manufacture Details

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining

software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon. In some cases, some aspects
5 of the invention may be implemented in hardware in an appliance that a migration team brings into a customer environment.

One or more embodiments of the invention, or elements thereof, can be implemented in the form of an apparatus including a memory and at least one processor that is coupled to the memory and operative to perform exemplary method steps.

10 One or more embodiments can make use of software running on a general purpose computer or workstation. With reference to FIG. 1, such an implementation might employ, for example, a processor 16, a memory 28, and an input/output interface 22 to a display 24 and external device(s) 14 such as a keyboard, a pointing device, or the like. The term “processor” as used herein is intended to include any processing device,
15 such as, for example, one that includes a CPU (central processing unit) and/or other forms of processing circuitry. Further, the term “processor” may refer to more than one individual processor. The term “memory” is intended to include memory associated with a processor or CPU, such as, for example, RAM (random access memory) 30, ROM (read only memory), a fixed memory device (for example, hard drive 34), a removable memory
20 device (for example, diskette), a flash memory and the like. In addition, the phrase “input/output interface” as used herein, is intended to contemplate an interface to, for example, one or more mechanisms for inputting data to the processing unit (for example, mouse), and one or more mechanisms for providing results associated with the processing unit (for example, printer). The processor 16, memory 28, and input/output interface 22
25 can be interconnected, for example, via bus 18 as part of a data processing unit 12. Suitable interconnections, for example via bus 18, can also be provided to a network interface 20, such as a network card, which can be provided to interface with a computer network, and to a media interface, such as a diskette or CD-ROM drive, which can be provided to interface with suitable media.

30 Accordingly, computer software including instructions or code for performing the methodologies of the invention, as described herein, may be stored in one or more of the

associated memory devices (for example, ROM, fixed or removable memory) and, when ready to be utilized, loaded in part or in whole (for example, into RAM) and implemented by a CPU. Such software could include, but is not limited to, firmware, resident software, microcode, and the like.

5 A data processing system suitable for storing and/or executing program code will include at least one processor 16 coupled directly or indirectly to memory elements 28 through a system bus 18. The memory elements can include local memory employed during actual implementation of the program code, bulk storage, and cache memories 32 which provide temporary storage of at least some program code in order to reduce the
10 number of times code must be retrieved from bulk storage during implementation.

Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, and the like) can be coupled to the system either directly or through intervening I/O controllers.

Network adapters 20 may also be coupled to the system to enable the data
15 processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

As used herein, including the claims, a “server” includes a physical data
20 processing system (for example, system 12 as shown in FIG. 1) running a server program. It will be understood that such a physical server may or may not include a display and keyboard.

As noted, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having
25 computer readable program code embodied thereon. Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any
30 suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical

connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any
5 suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable signal medium may include a propagated data signal with
10 computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a
15 program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

20 Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. In the most general case, the program code
25 may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to
30 an external computer (for example, through the Internet using an Internet Service Provider). However, one or more embodiments are particularly significant in the context

of a cloud or virtual machine environment. Reference is made back to FIGS. 1-3 and accompanying text.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some

alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each
5 block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

It should be noted that any of the methods described herein can include an
10 additional step of providing a system comprising distinct software modules embodied on a computer readable storage medium; the modules can include, for example, any or all of the appropriate elements depicted in the block diagrams in any of the pertinent figures and/or described herein; by way of example and not limitation, any one, some or all of the modules/blocks and or sub-modules/sub-blocks in FIGS. 16-19, such as a discovery
15 tool module to carry out functionality 16100 (optionally with sub-modules to carry out 16110, 16120, 16130, and/or 16140), a description module to hold description 16300 (optionally with sub-data-structures for data 16310, 16320, 16330, and/or 16340), and an infrastructure comparison engine module 16200 (optionally with sub-modules to implement functionality 16210, 16220, 16230, and/or 16240). The customer
20 environment 1702 may have one or more modules to implement one or more transfer agents and/or may have some or all of the discovery tool functionality as discussed elsewhere. The method steps can then be carried out using the distinct software modules and/or sub-modules of the system, as described above, executing on one or more hardware processors such as 16. Further, a computer program product can include a
25 computer-readable storage medium with code adapted to be implemented to carry out one or more method steps described herein, including the provision of the system with the distinct software modules.

In any case, it should be understood that the components illustrated herein may be implemented in various forms of hardware, software, or combinations thereof; for
30 example, application specific integrated circuit(s) (ASICs), functional circuitry, one or more appropriately programmed general purpose digital computers with associated

memory, and the like. Given the teachings of the invention provided herein, one of ordinary skill in the related art will be able to contemplate other implementations of the components of the invention.

5 The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not
10 preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as
15 specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the
20 principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

Claims

What is claimed is:

1. A non-transitory computer readable medium comprising computer executable instructions which when executed by a computer cause the computer to perform the method of:

discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component, wherein said at least one source infrastructure management component is an instance of an image, and wherein said at least one source infrastructure management component is running in a customer environment;

querying a database to obtain a description of a target cloud infrastructure;

analyzing said at least one source infrastructure management component using said description of said target cloud infrastructure to determine that said at least one source infrastructure management component is appropriate for infrastructure configuration mapping to said target cloud infrastructure;

stopping an application executing on said at least one source infrastructure management component determined appropriate for infrastructure configuration mapping; and

capturing said at least one source infrastructure management component determined appropriate for infrastructure configuration mapping for migration to said target cloud infrastructure.

2. The non-transitory computer readable medium of Claim 1, wherein, in said discovering step of said method, said at least one source infrastructure management component comprises at least one of:

at least one source infrastructure management client;

at least one source infrastructure management server;

at least one source infrastructure management configuration; and

at least one source infrastructure management log.

3. The non-transitory computer readable medium of Claim 2, wherein said computer executable instructions further cause the computer to perform the additional method step of deriving a description of non-functional requirements of said source management infrastructure.

5

4. The non-transitory computer readable medium of Claim 3, wherein:
in said discovering step of said method, said at least one source infrastructure management component comprises at least:

said at least one source infrastructure management configuration; and

10

said at least one source infrastructure management log;

in said querying step of said method, said description of said target cloud infrastructure comprises at least one of:

cloud infrastructure software standards;

cloud infrastructure software configurations; and

15

a description of non-functional requirements of a target management infrastructure of said target cloud infrastructure;

said analyzing step of said method comprises:

matching said cloud infrastructure software standards with at least a corresponding one of:

20

said at least one source infrastructure management client; and

said at least one source infrastructure management server;

mapping said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure; and

25

mapping said cloud infrastructure software configurations with said at least one source infrastructure management configuration and said at least one source infrastructure management log.

5. The non-transitory computer readable medium of Claim 4, wherein said mapping
30 of said cloud infrastructure software configurations with said at least one source

infrastructure management configuration and said at least one source infrastructure management log is at least partially based on:

said matching of said cloud infrastructure software standards with said at least one source infrastructure management client and said at least one source infrastructure management server; and

said mapping of said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure.

6. The non-transitory computer readable medium of Claim 5, wherein:

in said discovering step of said method, said at least one source infrastructure management component further comprises at least one source infrastructure management process;

in said querying step of said method, said description of said target cloud infrastructure further comprises at least one target infrastructure management process; and

said analyzing step of said method further comprises mapping said at least one source infrastructure management process with said at least one target infrastructure management process.

20

7. The non-transitory computer readable medium of Claim 6, wherein said mapping of said at least one source infrastructure management process with said at least one target infrastructure management process is at least partially based on said mapping of said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure.

8. The non-transitory computer readable medium of Claim 2, wherein:

in said querying step of said method, said description of said target cloud infrastructure comprises at least one of:

cloud infrastructure software standards; and

cloud infrastructure software configurations;
said analyzing step of said method comprises:

matching said cloud infrastructure software standards with at least a
corresponding one of:

5 said at least one source infrastructure management client; and
 said at least one source infrastructure management server; and
 mapping said cloud infrastructure software configurations with said at
least one source infrastructure management configuration.

10 9. The non-transitory computer readable medium of Claim 8, wherein said mapping
of said cloud infrastructure software configurations with said at least one source
infrastructure management configuration is at least partially based on said matching of
said cloud infrastructure software standards with said at least one source infrastructure
management client and said at least one source infrastructure management server.

15 10. The non-transitory computer readable medium of Claim 9, wherein:
 in said discovering step of said method, said at least one source infrastructure
management component further comprises at least one source infrastructure management
process;

20 in said querying step of said method, said description of said target cloud
infrastructure further comprises at least one target infrastructure management process;
and

 said analyzing step of said method further comprises mapping said at least one
source infrastructure management process with said at least one target infrastructure
25 management process.

11. The non-transitory computer readable medium of Claim 1, wherein said analyzing
step of said method further comprises determining that at least one conflict exists with
said at least one mandatory target infrastructure management component, wherein said at
30 least one conflict comprises said at least one source infrastructure management

component managing at least one object that said at least one mandatory target infrastructure management component will manage in said target cloud infrastructure.

12. The non-transitory computer readable medium of Claim 1, wherein said analyzing
5 step of said method further comprises determining that at least one conflict exists with said at least one mandatory target infrastructure management component, wherein said at least one conflict comprises said at least one source infrastructure management component using at least one resource that said at least one mandatory target infrastructure management component will use in said target cloud infrastructure.

10

13. The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method further comprises determining that at least one conflict exists with said at least one mandatory target infrastructure management component, wherein said at least one conflict comprises at least one of:

15 current lack of a mandatory target infrastructure component on a client,
existence of a mandatory target infrastructure component in a different version,
and
existence of a mandatory target infrastructure component with different configurations.

20

14. The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method comprises at least one of:

recommending un-installation of a source infrastructure component;
recommending installation of a target infrastructure component;
25 recommending modifying configuration of an infrastructure component; and
recommending excluding servers with a given source infrastructure component from migration.

15. The non-transitory computer readable medium of Claim 1, wherein:
30 said discovering step of said method comprises discovering, in said source computing system, a plurality of source infrastructure management components; and

said analyzing step of said method comprises:

initially aggregating all of said plurality of source infrastructure management components;

5 developing generic recommendations for said aggregated plurality of source infrastructure management components;

applying said recommendations on a per-instance basis; and

refining at least some of said applied recommendations in cases where said generic recommendations are inconclusive.

10 16. The non-transitory computer readable medium of Claim 1, wherein said discovering step of said method is carried out on at least one physical machine in said source computing system, said at least one physical machine comprising a client of said at least one source infrastructure management component.

15 17. The non-transitory computer readable medium of Claim 1, wherein said discovering step of said method is carried out on at least one physical machine in said source computing system, said at least one physical machine comprising a server on which said at least one source infrastructure management component runs.

20 18. The non-transitory computer readable medium of Claim 1, wherein said computer executable instructions comprise distinct software modules, each of the distinct software modules being embodied on said non-transitory computer readable medium, and wherein the distinct software modules comprise a discovery tool module, a description module, and an infrastructure comparison engine module;

25 wherein:

said discovering is carried out by said discovery tool module executing on said computer;

said querying is carried out by said description module executing on said computer; and

30 said analyzing is carried out by said infrastructure comparison engine module executing on said computer.

19. A non-transitory computer readable medium comprising computer executable instructions which when executed by a computer cause the computer to perform the method of:

5 discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component, wherein said at least one source infrastructure management component is an instance of an image, and wherein said at least one source infrastructure management component is running in a customer environment;

10 querying a database to obtain a description of a target cloud infrastructure;
 analyzing said at least one source infrastructure management component using said description of said target cloud infrastructure to determine that said at least one source infrastructure management component is appropriate for infrastructure configuration mapping to said target cloud infrastructure,

15 wherein:

 in said discovering step, said at least one source infrastructure management component comprises at least one of:

 at least one source infrastructure management client;
 at least one source infrastructure management server;
20 at least one source infrastructure management configuration; and
 at least one source infrastructure management log;

 in said querying step, said description of said target cloud infrastructure comprises at least one of:

 cloud infrastructure software standards; and
25 cloud infrastructure software configurations;

 said analyzing comprises:

 matching said cloud infrastructure software standards with at least a corresponding one of:

 said at least one source infrastructure management client; and
30 said at least one source infrastructure management server; and

mapping said cloud infrastructure software configurations with said at least one source infrastructure management configuration determined appropriate for infrastructure configuration mapping.

Abstract

In a source computing system having a source management infrastructure, at least one source infrastructure management component is discovered. A description of a target cloud infrastructure having a target management infrastructure is obtained. The description includes at least one mandatory target infrastructure management component. The at least one source infrastructure management component is analyzed to determine whether at least one conflict exists with the at least one mandatory target infrastructure management component.

10

15

20

Electronic Patent Application Fee Transmittal				
Application Number:				
Filing Date:				
Title of Invention:		MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		
First Named Inventor/Applicant Name:		Matthew A. Markley		
Filer:		Paul J. Otterstedt		
Attorney Docket Number:		YOR920120450US2		
Filed as Large Entity				
Filing Fees for Utility under 35 USC 111(a)				
Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Basic Filing:				
Utility application filing	1011	1	280	280
Utility Search Fee	1111	1	600	600
Utility Examination Fee	1311	1	720	720
Pages:				
Utility Appl Size fee per 50 sheets >100	1081	1	400	400
Claims:				
Miscellaneous-Filing:				
Petition:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
Total in USD (\$)				2000

Electronic Acknowledgement Receipt

EFS ID:	25104989
Application Number:	15061021
International Application Number:	
Confirmation Number:	4734
Title of Invention:	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
First Named Inventor/Applicant Name:	Matthew A. Markley
Customer Number:	98841
Filer:	Paul J. Otterstedt
Filer Authorized By:	
Attorney Docket Number:	YOR920120450US2
Receipt Date:	04-MAR-2016
Filing Date:	
Time Stamp:	13:52:38
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	\$2000
RAM confirmation Number	152
Deposit Account	500510
Authorized User	OTTERSTEDT, ELLENBOGEN & KAMMER, LLP

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

Charge any Additional Fees required under 37 CFR 1.16 (National application filing, search, and examination fees)

Charge any Additional Fees required under 37 CFR 1.17 (Patent application and reexamination processing fees)

Charge any Additional Fees required under 37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Power of Attorney	YOR920120450US2_PowerofAttorney_03-04-16.pdf	324576	no	2
			94c0f0e9e980c5755fe8d3fc371108f914f87c52		
Warnings:					
Information:					
2	Application Data Sheet	YOR920120450US2_ADS_03-04-16.pdf	1793477	no	9
			246bbec34f8fc03a7c3da113e1a4ea0751708e87		
Warnings:					
Information:					
3	Drawings-other than black and white line drawings	YOR920120450US2_FormalDrawings_prn.pdf	8556016	no	47
			acfbfd753df5581716c2912382977009df253736b		
Warnings:					
Information:					
4		YOR920120450US2_FiledApplication.pdf	438624	yes	120
			71fb2f45909ce53c8687f16267684ceb45cf5942		
	Multipart Description/PDF files in .zip description				
	Document Description		Start	End	
	Specification		1	111	
	Claims		112	119	
	Abstract		120	120	
Warnings:					
Information:					
5	Fee Worksheet (SB06)	fee-info.pdf	37047	no	2
			6cc563a5893faec93d06bd787501330d0e44d670		
Warnings:					
Information:					
Total Files Size (in bytes):			11149740		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT ASSIGNMENT COVER SHEETElectronic Version v1.1
Stylesheet Version v1.2

EPAS ID: PAT3768729

SUBMISSION TYPE:	NEW ASSIGNMENT
NATURE OF CONVEYANCE:	ASSIGNMENT
CONVEYING PARTY DATA	
Name	Execution Date
MATTHEW A. MARKLEY	11/28/2012
AMITKUMAR M. PARADKAR	11/26/2012
VENKATA VINAY PARISA	11/27/2012
BIRGIT M. PFITZMANN	11/12/2012
RECEIVING PARTY DATA	
Name:	INTERNATIONAL BUSINESS MACHINES CORPORATION
Street Address:	New Orchard Road
City:	Armonk
State/Country:	NEW YORK
Postal Code:	10504
PROPERTY NUMBERS Total: 1	
Property Type	Number
Application Number:	15061021
CORRESPONDENCE DATA	
Fax Number:	(203)769-5706
<i>Correspondence will be sent to the e-mail address first; if that is unsuccessful, it will be sent using a fax number, if provided; if that is unsuccessful, it will be sent via US Mail.</i>	
Phone:	203 433 3157
Email:	docket@oeklaw.com
Correspondent Name:	OTTERSTEDT, ELLENBOGEN & KAMMER, LLP
Address Line 1:	P.O. BOX 381
Address Line 4:	COS COB, CONNECTICUT 06807-0381
ATTORNEY DOCKET NUMBER:	YOR920120450US2
NAME OF SUBMITTER:	PAUL J. OTTERSTEDT
SIGNATURE:	/Paul J. Otterstedt/
DATE SIGNED:	03/04/2016
This document serves as an Oath/Declaration (37 CFR 1.63).	
Total Attachments: 8	
source=YOR920120450US2_FiledAssignment_03-04-16#page1.tif	

source=YOR920120450US2_FiledAssignment_03-04-16#page2.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page3.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page4.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page5.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page6.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page7.tif
source=YOR920120450US2_FiledAssignment_03-04-16#page8.tif

COPY

IBM DOCKET NUMBER: YOR920120450US1

**DECLARATION (37 C.F.R. 1.63) FOR UTILITY PATENT APPLICATION USING AN
APPLICATION DATA SHEET (37 C.F.R. 1.76) AND ASSIGNMENT**

Title of Invention: **MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD
MIGRATION**

As a below named inventor, I hereby declare that:

This declaration is directed to the attached application, or (if following box is checked):

☐ United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the
application.

I have reviewed and understand the contents of the application, including the claims.

I am aware of the duty to disclose to the United States Patent and Trademark Office all information
known to me to be material to patentability as defined in 37 CFR Section 1.56.

Whereas, the undersigned inventor(s) has/have made certain inventions, improvements, and discoveries
(herein referred to as the "Invention") disclosed in the above-identified patent application and further
identified by the IBM Docket Number provided above in the header of this document;

Whereas, International Business Machines Corporation, a corporation of New York having a place of
business at Armonk, New York (herein referred to as "IBM"), desires to acquire, and each undersigned
inventor desires to grant to IBM, the entire worldwide right, title, and interest in and to the Invention and
in and to any and all patent applications and patents directed thereto;

Now, therefore, for good and valuable consideration, the receipt and sufficiency thereof being hereby
acknowledged, each undersigned inventor ("ASSIGNOR") hereby sells or has sold, assigns or has
assigned, and otherwise transfers or has transferred to IBM (the "ASSIGNEE"), its successors, legal
representatives, and assigns, the entire worldwide right, title, and interest in and to the Invention, the
above-identified United States patent application, and any and all other patent applications and patents for
the Invention which may be applied for or granted therefor in the United States and in all foreign
countries and jurisdictions, including all divisions, continuations, reissues, reexaminations, renewals,
extensions, counterparts, substitutes, and extensions thereof, and all rights of priority resulting from the
filing of such applications and granting of such patents. In addition, each undersigned inventor hereby
authorizes and requests the Director of the United States Patent and Trademark Office to issue any United
States Patent, and foreign patent authorities to issue any foreign patent, granted for the Invention, to IBM,
its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the
same to be held and enjoyed by IBM, its successors, legal representatives, and assigns to the full end of
the terms for which any and all such patents may be granted, as fully and entirely as would have been
held and enjoyed by the undersigned had this Assignment not been made; and each undersigned inventor
agrees to execute any and all documents and instruments and perform all lawful acts reasonably related to
recording this Assignment or perfecting title to the Invention and all related patents and applications, in
IBM, its successors, legal representatives, and assigns, whenever requested by IBM, its successors, legal
representatives, or assigns.

COPY

IBM DOCKET NUMBER: YOR920120450US1

Each undersigned inventor acknowledges their prior and ongoing obligations to sell, assign, and transfer the rights under this Assignment to IBM and is unaware of any reason why they may not have the full and unencumbered right to sell, assign, and transfer the rights hereby sold, assigned, and transferred, and has not executed, and will not execute, any document or instrument in conflict herewith. Each undersigned inventor also hereby grants IBM, its successors, legal representatives, and assigns, the right to insert in this Assignment any further identification (including, but not limited to, patent Application Number) which may be necessary or desirable for recordation of this Assignment. This Assignment is governed by the substantive laws of the State of New York, and any disputes will be resolved in a New York state court or federal court sited in New York.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

(1) Legal Name of Inventor: Matthew A. Markley

Signature: Matthew A. Markley Date: 11-28-2012

(2) Legal Name of Inventor: Amitkumar M. Paradkar

Signature: _____ Date: _____

(3) Legal Name of Inventor: Venkata Vinay Parisa

Signature: _____ Date: _____

(4) Legal Name of Inventor: Birgit M. Pfitzmann

Signature: _____ Date: _____

COPY

IBM DOCKET NUMBER: YOR920120450US1

DECLARATION (37 C.F.R. 1.63) FOR UTILITY PATENT APPLICATION USING AN APPLICATION DATA SHEET (37 C.F.R. 1.76) AND ASSIGNMENT

Title of Invention: **MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION**

As a below named inventor, I hereby declare that:

This declaration is directed to the attached application, or (if following box is checked):

☐ United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I have reviewed and understand the contents of the application, including the claims.

I am aware of the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in 37 CFR Section 1.56.

Whereas, the undersigned inventor(s) has/have made certain inventions, improvements, and discoveries (herein referred to as the "Invention") disclosed in the above-identified patent application and further identified by the IBM Docket Number provided above in the header of this document;

Whereas, International Business Machines Corporation, a corporation of New York having a place of business at Armonk, New York (herein referred to as "IBM"), desires to acquire, and each undersigned inventor desires to grant to IBM, the entire worldwide right, title, and interest in and to the Invention and in and to any and all patent applications and patents directed thereto;

Now, therefore, for good and valuable consideration, the receipt and sufficiency thereof being hereby acknowledged, each undersigned inventor ("ASSIGNOR") hereby sells or has sold, assigns or has assigned, and otherwise transfers or has transferred to IBM (the "ASSIGNEE"), its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the Invention, the above-identified United States patent application, and any and all other patent applications and patents for the Invention which may be applied for or granted therefor in the United States and in all foreign countries and jurisdictions, including all divisions, continuations, reissues, reexaminations, renewals, extensions, counterparts, substitutes, and extensions thereof, and all rights of priority resulting from the filing of such applications and granting of such patents. In addition, each undersigned inventor hereby authorizes and requests the Director of the United States Patent and Trademark Office to issue any United States Patent, and foreign patent authorities to issue any foreign patent, granted for the Invention, to IBM, its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the same to be held and enjoyed by IBM, its successors, legal representatives, and assigns to the full end of the terms for which any and all such patents may be granted, as fully and entirely as would have been held and enjoyed by the undersigned had this Assignment not been made; and each undersigned inventor agrees to execute any and all documents and instruments and perform all lawful acts reasonably related to recording this Assignment or perfecting title to the Invention and all related patents and applications, in IBM, its successors, legal representatives, and assigns, whenever requested by IBM, its successors, legal representatives, or assigns.

COPY

IBM DOCKET NUMBER: YOR920120450UST

Each undersigned inventor acknowledges their prior and ongoing obligations to sell, assign, and transfer the rights under this Assignment to IBM and is unaware of any reason why they may not have the full and unencumbered right to sell, assign, and transfer the rights hereby sold, assigned, and transferred, and has not executed, and will not execute, any document or instrument in conflict herewith. Each undersigned inventor also hereby grants IBM, its successors, legal representatives, and assigns, the right to insert in this Assignment any further identification (including, but not limited to, patent Application Number) which may be necessary or desirable for recordation of this Assignment. This Assignment is governed by the substantive laws of the State of New York, and any disputes will be resolved in a New York state court or federal court sited in New York.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

(1) Legal Name of Inventor: Matthew A. Markley

Signature: _____ Date: _____

(2) Legal Name of Inventor: Amitkumar M. Paradkar

Signature: Amitkumar M. Paradkar Date: 11/26/2012

(3) Legal Name of Inventor: Venkatesh Vinay Parisa

Signature: _____ Date: _____

(4) Legal Name of Inventor: Birgit M. Pfitzmann

Signature: _____ Date: _____

COPY

IBM DOCKET NUMBER: YOR920120450US1

DECLARATION (37 C.F.R. 1.63) FOR UTILITY PATENT APPLICATION USING AN APPLICATION DATA SHEET (37 C.F.R. 1.76) AND ASSIGNMENT

Title of Invention: **MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION**

As a below named inventor, I hereby declare that:

This declaration is directed to the attached application, or (if following box is checked):

☐ United States application or PCT international application number _____
filed on _____

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I have reviewed and understand the contents of the application, including the claims.

I am aware of the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in 37 CFR Section 1.56.

Whereas, the undersigned inventor(s) has/has made certain inventions, improvements, and discoveries (herein referred to as the "Invention") disclosed in the above-identified patent application and further identified by the IBM Docket Number provided above in the header of this document;

Whereas, International Business Machines Corporation, a corporation of New York having a place of business at Armonk, New York (herein referred to as "IBM"), desires to acquire, and each undersigned inventor desires to grant to IBM, the entire worldwide right, title, and interest in and to the Invention and in and to any and all patent applications and patents directed thereto;

Now, therefore, for good and valuable consideration, the receipt and sufficiency thereof being hereby acknowledged, each undersigned inventor ("ASSIGNOR") hereby sells or has sold, assigns or has assigned, and otherwise transfers or has transferred to IBM (the "ASSIGNEE"), its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the Invention, the above-identified United States patent application, and any and all other patent applications and patents for the Invention which may be applied for or granted therefor in the United States and in all foreign countries and jurisdictions, including all divisions, continuations, reissues, reexaminations, renewals, extensions, counterparts, substitutes, and extensions thereof, and all rights of priority resulting from the filing of such applications and granting of such patents. In addition, each undersigned inventor hereby authorizes and requests the Director of the United States Patent and Trademark Office to issue any United States Patent, and foreign patent authorities to issue any foreign patent, granted for the Invention, to IBM, its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the same to be held and enjoyed by IBM, its successors, legal representatives, and assigns to the full end of the terms for which any and all such patents may be granted, as fully and entirely as would have been held and enjoyed by the undersigned had this Assignment not been made; and each undersigned inventor agrees to execute any and all documents and instruments and perform all lawful acts reasonably related to recording this Assignment or perfecting title to the Invention and all related patents and applications, in IBM, its successors, legal representatives, and assigns, whenever requested by IBM, its successors, legal representatives, or assigns.

COPY

IBM DOCKET NUMBER: YOR920120450US1

Each undersigned inventor acknowledges their prior and ongoing obligations to sell, assign, and transfer the rights under this Assignment to IBM and is unaware of any reason why they may not have the full and unencumbered right to sell, assign, and transfer the rights hereby sold, assigned, and transferred, and has not executed, and will not execute, any document or instrument in conflict herewith. Each undersigned inventor also hereby grants IBM, its successors, legal representatives, and assigns, the right to insert in this Assignment any further identification (including, but not limited to, patent Application Number) which may be necessary or desirable for recordation of this Assignment. This Assignment is governed by the substantive laws of the State of New York, and any disputes will be resolved in a New York state court or federal court sited in New York.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.


(1) Legal Name of Inventor: Matthew A. Markley

Signature: _____ Date: _____

(2) Legal Name of Inventor: Amitkumar M. Paradkar

Signature: _____ Date: _____

(3) Legal Name of Inventor: Venkata Vinay Parisa

Signature:  _____ Date: 11/27/2012

(4) Legal Name of Inventor: Birgit M. Pfitzmann

Signature: _____ Date: _____

COPY

IBM DOCKET NUMBER: YOR920120450US1

**DECLARATION (37 C.F.R. 1.63) FOR UTILITY PATENT APPLICATION USING AN
APPLICATION DATA SHEET (37 C.F.R. 1.76) AND ASSIGNMENT**

Title of Invention: **MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD
MIGRATION**

As a below named inventor, I hereby declare that:

This declaration is directed to the attached application, or (if following box is checked):

☐ United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I have reviewed and understand the contents of the application, including the claims.

I am aware of the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in 37 CFR Section 1.56.

Whereas, the undersigned inventor(s) has/have made certain inventions, improvements, and discoveries (herein referred to as the "invention") disclosed in the above-identified patent application and further identified by the IBM Docket Number provided above in the header of this document;

Whereas, International Business Machines Corporation, a corporation of New York having a place of business at Armonk, New York (herein referred to as "IBM"), desires to acquire, and each undersigned inventor desires to grant to IBM, the entire worldwide right, title, and interest in and to the invention and in and to any and all patent applications and patents directed thereto;

Now, therefore, for good and valuable consideration, the receipt and sufficiency thereof being hereby acknowledged, each undersigned inventor ("ASSIGNOR") hereby sells or has sold, assigns or has assigned, and otherwise transfers or has transferred to IBM (the "ASSIGNEE"), its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the invention, the above-identified United States patent application, and any and all other patent applications and patents for the invention which may be applied for or granted therefor in the United States and in all foreign countries and jurisdictions, including all divisions, continuations, reissues, reexaminations, renewals, extensions, counterparts, substitutes, and extensions thereof, and all rights of priority resulting from the filing of such applications and granting of such patents. In addition, each undersigned inventor hereby authorizes and requests the Director of the United States Patent and Trademark Office to issue any United States Patent, and foreign patent authorities to issue any foreign patent, granted for the invention, to IBM, its successors, legal representatives, and assigns, the entire worldwide right, title, and interest in and to the same to be held and enjoyed by IBM, its successors, legal representatives, and assigns to the full end of the terms for which any and all such patents may be granted, as fully and entirely as would have been held and enjoyed by the undersigned had this Assignment not been made; and each undersigned inventor agrees to execute any and all documents and instruments and perform all lawful acts reasonably related to recording this Assignment or perfecting title to the invention and all related patents and applications, in IBM, its successors, legal representatives, and assigns, whenever requested by IBM, its successors, legal representatives, or assigns.

COPY

IBM DOCKET NUMBER: YOR920120450US1

Each undersigned inventor acknowledges their prior and ongoing obligations to sell, assign, and transfer the rights under this Assignment to IBM and is unaware of any reason why they may not have the full and unencumbered right to sell, assign, and transfer the rights hereby sold, assigned, and transferred, and has not executed, and will not execute, any document or instrument in conflict herewith. Each undersigned inventor also hereby grants IBM, its successors, legal representatives, and assigns, the right to insert in this Assignment any further identification (including, but not limited to, patent Application Number) which may be necessary or desirable for recordation of this Assignment. This Assignment is governed by the substantive laws of the State of New York, and any disputes will be resolved in a New York state court or federal court sited in New York.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

(1) Legal Name of Inventor: Matthew A. Markley

Signature: _____ Date: _____

(2) Legal Name of Inventor: Amitkumar M. Paradkar

Signature: _____ Date: _____

(3) Legal Name of Inventor: Venkata Vinay Parisa

Signature: _____ Date: _____

(4) Legal Name of Inventor: Birgit M. Pfitzmann

Signature: Birgit Pfitzmann Date: 11/12/2012



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY. DOCKET NO	TOT CLAIMS	IND CLAIMS
15/061,021	03/04/2016	2447	2000	YOR920120450US2	19	2

CONFIRMATION NO. 4734

98841

Otterstedt, Ellenbogen & Kammer, LLP
 P.O. Box 381
 Cos Cob, CT 06807-0381

FILING RECEIPT



0000000081533888

Date Mailed: 03/21/2016

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Matthew A. Markley, Hays, KS;
 Amitkumar M. Paradkar, Mohegan Lake, NY;
 Venkata Vinay Parisa, Charlotte, NC;
 Birgit M. Pfitzmann, Samstagern, SWITZERLAND;

Applicant(s)

International Business Machines Corporation, Armonk, NY;

Power of Attorney: The patent practitioners associated with Customer Number 98841

Domestic Priority data as claimed by applicant

This application is a CON of 13/689,500 11/29/2012 PAT 9282166

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see <http://www.uspto.gov> for more information.) - None.

Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: Yes

Permission to Access Search Results: Yes

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 03/18/2016

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 15/061,021**

Projected Publication Date: 06/30/2016

Non-Publication Request: No

Early Publication Request: No

Title

MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

Preliminary Class

709

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at <http://www.uspto.gov/web/offices/pac/doc/general/index.html>.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, <http://www.stopfakes.gov>. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

LICENSE FOR FOREIGN FILING UNDER**Title 35, United States Code, Section 184****Title 37, Code of Federal Regulations, 5.11 & 5.15****GRANTED**

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign Assets Control, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit <http://www.SelectUSA.gov> or call +1-202-482-6800.

6361

PATENT APPLICATION FEE DETERMINATION RECORD						Application or Docket Number 15/061,021				
Substitute for Form PTO-875										
APPLICATION AS FILED - PART I										
(Column 1)		(Column 2)		SMALL ENTITY		OR OTHER THAN SMALL ENTITY				
FOR	NUMBER FILED	NUMBER EXTRA	RATE(\$)	FEE(\$)		RATE(\$)	FEE(\$)			
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	280			
SEARCH FEE (37 CFR 1.16(k), (l), or (m))	N/A	N/A	N/A			N/A	600			
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	720			
TOTAL CLAIMS (37 CFR 1.16(j))	19	minus 20 = *			OR	x 80 =	0.00			
INDEPENDENT CLAIMS (37 CFR 1.16(h))	2	minus 3 = *				x 420 =	0.00			
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is \$310 (\$155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						400			
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							0.00			
			TOTAL			TOTAL	2000			
* If the difference in column 1 is less than zero, enter "0" in column 2.										
APPLICATION AS AMENDED - PART II										
(Column 1)		(Column 2)		(Column 3)		SMALL ENTITY		OR OTHER THAN SMALL ENTITY		
AMENDMENT A		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE(\$)	ADDITIONAL FEE(\$)		RATE(\$)	ADDITIONAL FEE(\$)
	Total (37 CFR 1.16(i))	*	Minus	**	=	x	=	OR	x	=
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x	=	OR	x	=
	Application Size Fee (37 CFR 1.16(s))							OR		
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	
(Column 1)		(Column 2)		(Column 3)		SMALL ENTITY		OR OTHER THAN SMALL ENTITY		
AMENDMENT B		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE(\$)	ADDITIONAL FEE(\$)		RATE(\$)	ADDITIONAL FEE(\$)
	Total (37 CFR 1.16(i))	*	Minus	**	=	x	=	OR	x	=
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x	=	OR	x	=
	Application Size Fee (37 CFR 1.16(s))							OR		
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	
* If the entry in column 1 is less than the entry in column 2, write "0" in column 3. ** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20". *** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3". The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.										



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/061,021	03/04/2016	Matthew A. Markley	YOR920120450US2

CONFIRMATION NO. 4734

98841

Otterstedt, Ellenbogen & Kammer, LLP

P.O. Box 381

Cos Cob, CT 06807-0381

PUBLICATION NOTICE



OC00000084072537

Title:MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION**Publication No.**US-2016-0191298-A1**Publication Date:**06/30/2016**NOTICE OF PUBLICATION OF APPLICATION**

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently <http://www.uspto.gov/patft/>.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Office of Public Records. The Office of Public Records can be reached by telephone at (703) 308-9726 or (800) 972-6382, by facsimile at (703) 305-8759, by mail addressed to the United States Patent and Trademark Office, Office of Public Records, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently <http://pair.uspto.gov/>. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

Doc code: IDS

6363

PTO/SB/08a (03-15)

Doc description: Information Disclosure Statement (IDS) Filed

Approved for use through 07/31/2016. OMB 0651-0031

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

U.S.PATENTS

Remove

Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7844853		2010-11-30	Barsness et al.	
	2	7191441		2007-03-13	Abbott et al.	
	3	7356679		2008-04-08	Le et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS

Remove

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20120131567		2012-05-24	Barros et al.	
	2	20110320882		2011-12-29	Beaty et al.	
	3	20110107331		2011-05-05	Evans et al.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

4	20090307166	2009-12-10	Routray et al.
5	20090063751	2009-03-05	Dow, Eli M.
6	20100077160	2010-03-25	Liu et al.
7	20120117558	2012-05-10	Futty et al.
8	20140129521	2014-05-08	Marsden, Luke
9	20090288084	2009-11-19	Astete et al.
10	20130036091	2013-02-07	PROVENZANO et al.
11	20070033356	2007-02-08	Erlikhman, Boris
12	20140115285	2014-04-24	Arcese et al.
13	20140149695	2014-05-29	Zaslavsky et al.
14	20120266169	2012-10-18	Lagergren et al.

Application Number
6365

15061021

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

15	20130167145	2013-06-27	Krishnamurthy et al.
16	20100049929	2010-02-25	Nagarkar et al.
17	20120209812	2012-08-16	Bezbaruah et al.
18	20100191930	2010-07-29	Groff; Dana ; et al.
19	20140149983	2014-05-29	Milton A. Bonilla
20	20150331708	2015-11-29	Vasanth Bala

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS

Remove

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button

Add

NON-PATENT LITERATURE DOCUMENTS

Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵

Application Number
6366

15061021

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

1

If you wish to add additional non-patent literature document citation information please click the Add button

Add

EXAMINER SIGNATURE

Examiner Signature

Date Considered

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number
6367

15061021

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.**SIGNATURE**

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J. Otterstedt/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	Paul J. Otterstedt	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc code: IDS

6369

PTO/SB/08a (01-10)

Doc description: Information Disclosure Statement (IDS) Filed

Approved for use through 07/31/2012. OMB 0651-0031

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

U.S.PATENTS

Remove

Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS

Remove

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20120072903		2012-03-22	JOUKOV; NIKOLAI A. ; et al.	
	2	20120304179		2012-11-29	DEVARAKONDA; MURTHY ; et al.	
	3	20130091285		2013-04-11	Devarakonda; Murthy V. ; et al.	
	4	20060218544		2006-09-28	Chakraborty; Parag ; et al.	
	5	20110060832		2011-03-10	GOVIL; Kinshuk ; et al.	

Application Number
6370

15061021

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

6	20110202714	2011-08-18	Burch; Lloyd Leon ; et al.
7	20100287280	2010-11-11	Sivan; Gal
8	20100306772	2010-12-02	ARNOLD; WILLIAM C. ; et al.
9	20140149492	2014-05-29	Ananthanarayanan; Rema ; et al.
10	20140149591	2014-05-29	Bhattacharya; Kamal ; et al.
11	20140149494	2014-05-29	Markley; Matthew A. ; et al.
12	20140146055	2014-05-29	Bala; Vasanth ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS

Remove

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button

Add

NON-PATENT LITERATURE DOCUMENTS

Remove

Application Number
6371

15061021

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Chris Ward, et al. "Workload Migration into Clouds -- Challenges, Experiences, Opportunities". IEEE 3rd Intern. Conf. on Cloud Computing, 2010, pp. 164-171	
	2	Galapagos. "Model-driven discovery of end-to-end application-storage relationships in distributed systems". IBM Journal of Research and Development archive, Volume 52 Issue 4, July 2008, Pages 367-377	
	3	Nikolai Joukov, et al. "Application-Storage Discovery". SYSTOR 2010, Haifa, May 2010	
	4	Peter Mell, et al. "The NIST Definition of Cloud Computing". Version 15, 10-7-09, National Institute of Standards and Technology, Information Technology Laboratory	
	5	IBM Global Technology Services Data Sheet "Managed hosted cloud infrastructure services". IBM Corporation 2012	
	6	IBM Migration Factory, downloaded from http://www-03.ibm.com/systems/migratetoibm/factory/ on 11-12-2012	
	7	Robert Filepp et al., Image selection as a service for cloud computing environments; IEEE International Conference on Service-Oriented Computing and Applications (SOCA) 2010, 1-8 (abstract only)	

If you wish to add additional non-patent literature document citation information please click the Add button

Add

EXAMINER SIGNATURE

Examiner Signature

Date Considered

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number # 6372	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/PAUL J. OTTERSTEDT/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	PAUL J. OTTERSTEDT	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc code: IDS

#: 6374

PTO/SB/08a (01-10)

Doc description: Information Disclosure Statement (IDS) Filed

Approved for use through 07/31/2012. OMB 0651-0031

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

U.S.PATENTS

Remove

Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9282166		2016-03-08	Markley , et al.	
	2	9092837		2015-07-28	Bala , et al.	
	3	9292330		2016-03-22	Bonilla , et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS

Remove

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20050204040		2005-09-15	Ferri, Richard C. ; et al.	
	2	20120203823		2012-08-09	Gaurav Manglik	
	3	20120311154		2012-12-06	Christopher Edwin Morgan	

Application Number
6375

15061021

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

4	20120137003	2012-05-31	James Michael Ferris
5	20080189697	2008-08-07	Kachroo; Naveen K. ; et al.
6	20100235831	2010-09-16	Dittmer; Arend Erich
7	20110106768	2011-05-05	Khanzode; Pankaj V. ; et al.
8	20090249284	2009-10-01	Antosz; Wacław T. ; et al.
9	20160110215	2016-04-21	Bonilla; Milton A. ; et al.
10	20150331708	2015-11-19	Bala; Vasanth ; et al.
11	20160191332	2016-06-30	Markley; Matthew A. ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS

Remove

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button

Add

Application Number
6376

15061021

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)**NON-PATENT LITERATURE DOCUMENTS****Remove**

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1		

If you wish to add additional non-patent literature document citation information please click the Add button

Add**EXAMINER SIGNATURE**

Examiner Signature

Date Considered

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number # 6377	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J.Otterstedt/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	Paul J.Otterstedt	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	27603981
Application Number:	15061021
International Application Number:	
Confirmation Number:	4734
Title of Invention:	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
First Named Inventor/Applicant Name:	Matthew A. Markley
Customer Number:	98841
Filer:	Paul J. Otterstedt
Filer Authorized By:	
Attorney Docket Number:	YOR920120450US2
Receipt Date:	23-NOV-2016
Filing Date:	04-MAR-2016
Time Stamp:	14:11:12
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	YOR920120450US2_IDS_11-23-16_1of3.pdf	613149 b82ea94f9811616f351a2aeb6031d9e2674afcd9	no	5

Warnings:

Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	YOR920120450US2_IDS_11-23-16_2of3.pdf	1035867 fb8a88e3fd2dd85422e6ec988bd52b74d58b803b	no	6
Warnings:					
Information:					
3	Information Disclosure Statement (IDS) Form (SB08)	YOR920120450US2_IDS_11-23-16_3of3.pdf	612707 68ade45a35309da96d6e43d1e13af3874266ef2d	no	5
Warnings:					
Information:					
4	Non Patent Literature	YOR920120450US2_NPL1_11-23-16.pdf	1843502 cdef2b4ab1cae87124ac0f5ea65b914320e0c7e	no	8
Warnings:					
Information:					
5	Non Patent Literature	YOR920120450US2_NPL2_11-23-16.pdf	414682 aa72cf6ee3e16abed9a9df36962aca96f068a287	no	11
Warnings:					
Information:					
6	Non Patent Literature	YOR920120450US2_NPL3_11-23-16.pdf	2559949 929b597770813b4db8a74eec44133a8090702ac3	no	14
Warnings:					
Information:					
7	Non Patent Literature	YOR920120450US2_NPL4_11-23-16.pdf	29729 1c98e052bf6828fb7af2dbceffdb1fe4c0379f9f	no	2
Warnings:					
Information:					
8	Non Patent Literature	YOR920120450US2_NPL5_11-23-16.pdf	56038 464be1058766907cddf3f48ba68704a604a2a484	no	2
Warnings:					
Information:					

9	Non Patent Literature	YOR920120450US2_NPL6_11-23-16.pdf	82377 96d1e9d02ec31e9f386f1ec92911f6f38c9bbf56	no	4
---	-----------------------	-----------------------------------	---------------------------------------------------	----	---

Warnings:**Information:**

10	Non Patent Literature	YOR920120450US2_NPL7_11-23-16.pdf	142345 976661a4e1d1912c34c8e9fbaacb6a8b117ee2bc	no	2
----	-----------------------	-----------------------------------	----------------------------------------------------	----	---

Warnings:**Information:**

Total Files Size (in bytes):			7390345
-------------------------------------	--	--	---------

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/061,021	03/04/2016	Matthew A. Markley	YOR920120450US2	4734

98841 7590 12/16/2016
Otterstedt, Ellenbogen & Kammer, LLP
P.O. Box 381
Cos Cob, CT 06807-0381

EXAMINER

NGUYEN, TU T

ART UNIT	PAPER NUMBER
----------	--------------

2453

MAIL DATE	DELIVERY MODE
-----------	---------------

12/16/2016

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Application No.

15/061,021

Applicant(s)

MARKLEY ET AL.

Office Action Summary

Examiner

Tu Nguyen

Art Unit

2453

AIA (First Inventor to File)

Status

No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --**Period for Reply**

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☐ Responsive to communication(s) filed on _____.
☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____.; the restriction requirement and election have been incorporated into this action.
- 4) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

- 5) ☒ Claim(s) 1-19 is/are pending in the application.
 5a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 6) ☐ Claim(s) _____ is/are allowed.
- 7) ☒ Claim(s) 1-19 is/are rejected.
- 8) ☐ Claim(s) _____ is/are objected to.
- 9) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.

Application Papers

- 10) ☐ The specification is objected to by the Examiner.
- 11) ☒ The drawing(s) filed on 03/04/2016 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

- a) ☐ All b) ☐ Some** c) ☐ None of the:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☒ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b)
 Paper No(s)/Mail Date _____
- 3) ☐ Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____
- 4) ☐ Other: _____

Application/Control Number: 15/061,021
Art Unit: 2453

Page 2

The present application is being examined under the pre-AIA first to invent provisions.

DETAILED ACTION

Double Patenting

The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the “right to exclude” granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on nonstatutory double patenting provided the reference application or patent either is shown to be commonly owned with the examined application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement. See MPEP § 717.02 for applications subject to examination under the first inventor to file

Application/Control Number: 15/061,021
Art Unit: 2453

Page 3

provisions of the AIA as explained in MPEP § 2159. See MPEP §§ 706.02(l)(1) - 706.02(l)(3) for applications not subject to examination under the first inventor to file provisions of the AIA. A terminal disclaimer must be signed in compliance with 37 CFR 1.321(b).

The USPTO Internet website contains terminal disclaimer forms which may be used. Please visit www.uspto.gov/patent/patents-forms. The filing date of the application in which the form is filed determines what form (e.g., PTO/SB/25, PTO/SB/26, PTO/AIA/25, or PTO/AIA/26) should be used. A web-based eTerminal Disclaimer may be filled out completely online using web-screens. An eTerminal Disclaimer that meets all requirements is auto-processed and approved immediately upon submission. For more information about eTerminal Disclaimers, refer to www.uspto.gov/patents/process/file/efs/guidance/eTD-info-I.jsp.

Claims 1-19 are rejected on the ground of nonstatutory double patenting as being unpatentable over claims 1-19 of U.S. Patent No. 9,282,166 ('166 hereinafter). Although the claims at issue are not identical, they are not patentably distinct from each other because claims 1-19 of the current application are anticipated by claims 1-19 of '166.

With respect to claims 1-19, claims 1-19 of '166 discloses all the claimed limitations except for a non-transitory readable medium. However, the claimed medium would have been known. It would have been obvious to modify '166 by creating instructions for the method of claims 1-19 of '166 and store the instructions to the medium as claimed to use a computer to perform the method.

Application/Control Number: 15/061,021
Art Unit: 2453

Page 4

Allowable Subject Matter

Claims 1-19 would be allowable if applicant file in a Terminal Disclaimer to overcome the Double Patenting rejection, set forth in this Office action.

The prior art of record, alone or in any combination, does not teach or suggest the invention recited in independent claims 1, 19 since, among other things, the prior art of record does not teach or suggest claimed features relating to a non-transitory computer readable medium comprising computer executable instructions which when executed by a computer cause the computer to perform the method of: discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component, wherein said at least one source infrastructure management component is an instance of an image, and wherein said at least one source infrastructure management component is running in a customer environment; querying a database to obtain a description of a target cloud infrastructure; analyzing said at least one source infrastructure management component using said description of said target cloud infrastructure to determine that said at least one source infrastructure management component is appropriate for infrastructure configuration mapping to said target cloud infrastructure; stopping an application executing on said at least one source infrastructure management component determined appropriate for infrastructure configuration mapping, as recited in the context of claim 1, and similarly claimed in independent claim 19.

Application/Control Number: 15/061,021
Art Unit: 2453

Page 5

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Tu Nguyen whose telephone number is (571)272-2424. The examiner can normally be reached on M-W,F 7:30-5:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kamal Divecha can be reached on (571) 272-5863. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Tu Nguyen/
Primary Examiner, Art Unit 2453
12/11/2016

6388

Notice of References Cited	Application/Control No. 15/061,021		Applicant(s)/Patent Under Reexamination MARKLEY ET AL.	
	Examiner Tu Nguyen		Art Unit 2453	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-2012/0297441 A1	11-2012	Boldyrev; Sergey	G06F21/6245	726/1
*	B	US-2014/0059209 A1	02-2014	Alnoor; Ahmadullah	G06F9/54	709/224
	C	US-				
	D	US-				
	E	US-				
	F	US-				
	G	US-				
	H	US-				
	I	US-				
	J	US-				
	K	US-				
	L	US-				
	M	US-				

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		15061021
	Filing Date		2016-03-04
	First Named Inventor	Matthew A. Markley	
	Art Unit	2453	
	Examiner Name	NGUYEN, TU T	
	Attorney Docket Number	YOR920120450US2	

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9282166		2016-03-08	Markley , et al.	
	2	9092837		2015-07-28	Bala , et al.	
	3	9292330		2016-03-22	Bonilla , et al.	

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20050204040		2005-09-15	Ferri, Richard C. ; et al.	
	2	20120203823		2012-08-09	Gaurav Manglik	
	3	20120311154		2012-12-06	Christopher Edwin Morgan	

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

4	20120137003	2012-05-31	James Michael Ferris
5	20080189697	2008-08-07	Kachroo; Naveen K. ; et al.
6	20100235831	2010-09-16	Dittmer; Arend Erich
7	20110106768	2011-05-05	Khanzode; Pankaj V. ; et al.
8	20090249284	2009-10-01	Antosz; Wacław T. ; et al.
9	20160110215	2016-04-21	Bonilla; Milton A. ; et al.
10	20150331708	2015-11-19	Bala; Vasanth ; et al.
11	20160191332	2016-06-30	Markley; Matthew A. ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

NON-PATENT LITERATURE DOCUMENTS			Remove
Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1		
If you wish to add additional non-patent literature document citation information please click the Add button Add			
EXAMINER SIGNATURE			
Examiner Signature	/TU T NGUYEN/		Date Considered 12/11/2016
*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.			
<small>¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.</small>			

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.
The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J.Otterstedt/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	Paul J.Otterstedt	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc code: IDS

Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	15061021
	Filing Date	2016-03-04
	First Named Inventor	Matthew A. Markley
	Art Unit	2453
	Examiner Name	NGUYEN, TU T
	Attorney Docket Number	YOR920120450US2

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7844853		2010-11-30	Barsness et al.	
	2	7191441		2007-03-13	Abbott et al.	
	3	7356679		2008-04-08	Le et al.	
If you wish to add additional U.S. Patent citation information please click the Add button.						Add
U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20120131567		2012-05-24	Barros et al.	
	2	20110320882		2011-12-29	Beaty et al.	
	3	20110107331		2011-05-05	Evans et al.	

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

4	20090307166	2009-12-10	Routray et al.
5	20090063751	2009-03-05	Dow, Eli M.
6	20100077160	2010-03-25	Liu et al.
7	20120117558	2012-05-10	Futty et al.
8	20140129521	2014-05-08	Marsden, Luke
9	20090288084	2009-11-19	Astete et al.
10	20130036091	2013-02-07	PROVENZANO et al.
11	20070033356	2007-02-08	Erlikhman, Boris
12	20140115285	2014-04-24	Arcese et al.
13	20140149695	2014-05-29	Zaslavsky et al.
14	20120266169	2012-10-18	Lagergren et al.

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

15	20130167145	2013-06-27	Krishnamurthy et al.
16	20100049929	2010-02-25	Nagarkar et al.
17	20120209812	2012-08-16	Bezbaruah et al.
18	20100191930	2010-07-29	Groff; Dana ; et al.
19	20140149983	2014-05-29	Milton A. Bonilla
20	20150331708	2015-11-29	Vasanth Bala

If you wish to add additional U.S. Published Application citation information please click the Add button. [Add](#)

FOREIGN PATENT DOCUMENTS [Remove](#)

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button. [Add](#)

NON-PATENT LITERATURE DOCUMENTS [Remove](#)

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
--------------------	---------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

1	
---	--

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/TU T NGUYEN/	Date Considered	12/11/2016
--------------------	---------------	-----------------	------------

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.
The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J. Otterstedt/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	Paul J. Otterstedt	Registration Number	37411


This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Search Notes 	Application/Control No. 15061021	Applicant(s)/Patent Under Reexamination MARKLEY ET AL.
	Examiner TU NGUYEN	Art Unit 2453

CPC- SEARCHED		
Symbol	Date	Examiner

CPC COMBINATION SETS - SEARCHED		
Symbol	Date	Examiner

US CLASSIFICATION SEARCHED			
Class	Subclass	Date	Examiner

SEARCH NOTES		
Search Notes	Date	Examiner
East	12/11/2016	TN
709/203, 218-226	12/11/2016	TN
G06F 15/16, 15/173, 15/177, 9/455	12/11/2016	TN
H04L 29/08, 67/42, 29/06	12/11/2016	TN

INTERFERENCE SEARCH			
US Class/ CPC Symbol	US Subclass / CPC Group	Date	Examiner

--	--

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	48	("20060218544" "20100287280" "20100306772" "20110060832" "20110202714" "20120072903" "20120304179" "20130091285" "20140146055" "20140149492" "20140149494" "20140149591" "7844853" "7191441" "7356679" "20120131567" "20110320882" "20110107331" "20090307166" "20090063751" "20100077160" "20120117558" "20140129521" "20090288084" "20130036091" "20070033356" "20140115285" "20140149695" "20120266169" "20130167145" "20100049929" "20120209812" "20100191930" "20140149983" "20150331708" "9282166" "9092837" "9292330" "20050204040" "20120203823" "20120311154" "20120137003" "20080189697" "20100235831" "20110106768" "20090249284" "20160110215" "20160191332").PN.	US-PGPUB; USPAT	OR	OFF	2016/12/11 16:23
L4	20	((("7844853") or ("7191441") or ("7356679") or ("20120131567") or ("20110320882") or ("20110107331") or ("20090307166") or ("20090063751") or ("20100077160") or ("20120117558") or ("20140129521") or ("20090288084") or ("20130036091") or ("20070033356") or ("20140115285") or ("20140149695") or ("20120266169") or ("20130167145") or ("20100049929") or ("20120209812"))).PN.	US-PGPUB; USPAT	OR	OFF	2016/12/11 16:26
L5	13821	cloud with ((management or manager) near3 (component or module or system or apparatus))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L6	290	target with cloud with infra\$1structure	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L7	44	L5 and L6 and @ad<="20122911"	US-PGPUB; USPAT;	ADJ	ON	2016/12/11 16:26

6402

			USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			
L8	154570	(map or mapping) near7 \$5function\$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L9	6113326	configuration	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L10	315	L5 and L8 and L9 and @ad<="20122911"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L11	23338	cloud with infra\$1structure	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L12	238	L10 and L11	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L13	3332203	target	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L14	119	L12 and L13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L15	111479	query\$5 near3 database	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26

L16	24	L14 and L15	#: 6403 US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:26
L17	29	L4 or L7 or L16	USPAT	OR	OFF	2016/12/11 16:26
L18	3	(1 or 4) and 6	USPAT	OR	OFF	2016/12/11 16:28
L19	4260	cloud with migrat\$5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:37
L21	179	19 and 8 and 11 and @ad<="20121129"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:45
L22	13	21 and 15	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 16:49
L23	519	((program or application) near10 (run or running or execut\$5)) with ((map or mapping) near7 configuration)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:31
L24	1	21 and 23	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:33
L25	119	L12 and L13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:34
L26	134956	(709/203,218-226).OCLS.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/12/11 17:34

L27	17768	G06F15/16.cpc. or G06F15/173.cpc. or G06F15/177.cpc. or G06F9/455.cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:34
L28	99018	H04L29/08.cpc. or H04L67/42.cpc. or H04L29/06.cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:34
L29	236035	L26 or L27 or L28	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:34
L30	41	L29 and L25	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2016/12/11 17:34

12/ 11/ 2016 5:36:19 PM

C:\Users\nguyen24\Documents\EAST\Workspaces\15061021.wsp

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
 (Not for submission under 37 CFR 1.99)

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

U.S.PATENTS							Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	
	1						
If you wish to add additional U.S. Patent citation information please click the Add button.							Add
U.S.PATENT APPLICATION PUBLICATIONS							Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	
	1	20120072903		2012-03-22	JOUKOV; NIKOLAI A. ; et al.		
	2	20120304179		2012-11-29	DEVARAKONDA; MURTHY ; et al.		
	3	20130091285		2013-04-11	Devarakonda; Murthy V. ; et al.		
	4	20060218544		2006-09-28	Chakraborty; Parag ; et al.		
	5	20110060832		2011-03-10	GOVIL; Kinshuk ; et al.		

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

6	20110202714	2011-08-18	Burch; Lloyd Leon ; et al.
7	20100287280	2010-11-11	Sivan; Gal
8	20100306772	2010-12-02	ARNOLD; WILLIAM C. ; et al.
9	20140149492	2014-05-29	Ananthanarayanan; Rema ; et al.
10	20140149591	2014-05-29	Bhattacharya; Kamal ; et al.
11	20140149494	2014-05-29	Markley; Matthew A. ; et al.
12	20140146055	2014-05-29	Bala; Vasanth ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button. [Add](#)

FOREIGN PATENT DOCUMENTS [Remove](#)

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button. [Add](#)

NON-PATENT LITERATURE DOCUMENTS [Remove](#)

Application Number		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Chris Ward, et al. "Workload Migration into Clouds -- Challenges, Experiences, Opportunities". IEEE 3rd Intern. Conf. on Cloud Computing, 2010, pp. 164-171	
	2	Galapagos. "Model-driven discovery of end-to-end application-storage relationships in distributed systems". IBM Journal of Research and Development archive, Volume 52 Issue 4, July 2008, Pages 367-377	
	3	Nikolai Joukov, et al. "Application-Storage Discovery". SYSTOR 2010, Haifa, May 2010	
	4	Peter Mell, et al. " The NIST Definition of Cloud Computing". Version 15, 10-7-09, National Institute of Standards and Technology, Information Technology Laboratory	
	5	IBM Global Technology Services Data Sheet "Managed hosted cloud infrastructure services". IBM Corporation 2012	
	6	IBM Migration Factory, downloaded from http://www-03.ibm.com/systems/migratetoibm/factory/ on 11-12-2012	
	7	Robert Filepp et al., Image selection as a service for cloud computing environments; IEEE International Conference on Service-Oriented Computing and Applications (SOCA) 2010,1-8 (abstract only)	

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/TU T NGUYEN/	Date Considered	12/11/2016
--------------------	---------------	-----------------	------------

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number	15061021
Filing Date	2016-03-04
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.
The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/PAUL J. OTTERSTEDT/	Date (YYYY-MM-DD)	2016-11-23
Name/Print	PAUL J. OTTERSTEDT	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc code: IDS

6410

PTO/SB/08a (01-10)

Doc description: Information Disclosure Statement (IDS) Filed

Approved for use through 07/31/2012. OMB 0651-0031

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		15061021
	Filing Date		2016-03-04
	First Named Inventor	Matthew A. Markley	
	Art Unit	2453	
	Examiner Name	NGUYEN, TU T	
	Attorney Docket Number	YOR920120450US2	

U.S.PATENTS							Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	
	1	8566824		2013-10-22	Traut; Eric P.		
	2	8689054		2014-04-01	van der Goot; Roel		
	3	8726083		2014-05-13	van der Goot; Roel		
If you wish to add additional U.S. Patent citation information please click the Add button.							Add
U.S.PATENT APPLICATION PUBLICATIONS							Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	
	1	20140089658		2014-03-27	Raghuram; Yeluri ; et al.		
	2	20110295999		2011-12-01	Ferris; James Michael ; et al.		
	3	20130275596		2013-10-17	Subramaniam; Karthikeyan		

Application Number
0411

15061021

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

4	20140059539	2014-02-27	Simonsen; Harold C. ; et al.
5	20110071983	2011-03-24	Murase; Atsushi
6	20130262390	2013-10-03	Kumarasamy; Paramasivam ; et al.
7	20120084775	2012-04-05	Lotlikar; Mahesh ; et al.
8	20100106885	2010-04-29	Gao; Zhi Guo ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS

Remove

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	CN102387197	CN		2012-03-21	Lotlikar; Mahesh ; et al.		

If you wish to add additional Foreign Patent Document citation information please click the Add button

Add

NON-PATENT LITERATURE DOCUMENTS

Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Clark et al., "Live Migration of Virtual Machines," 2nd Symposium on Networked Systems Design & Implementation, NSDI'05, 2005, pp. 273-286.	

Application Number
6412

15061021

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

2	Nicolae, Bogdan and Cappello, Franck, "BlobCR: Efficient Checkpoint-Restart for HPC Applications on IaaS Clouds using Virtual Disk Snapshots," Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC'11, 12-18 November 2011, p. 1-12.
3	Park et al., "Fast and Space-Efficient Virtual Machine Checkpointing," Proceedings of the 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE'11, 9-11 March 2011, pp. 75-86.
4	Zhao, Ming and Figueiredo, Renato J., "Experimental Study of Virtual Machine Migration in Support of Reservation of Cluster Resources," Proceedings of the 2nd International Workshop on Virtualization Technology in Distributed Computing, VTDC'07, 12 November 2007, p 1-8.
5	state intellectual property office of the people's republic of China, First Office Action, Chinese Application 201310626100.0, June 28, 2016, includes machine translation, pages 1-8
6	state intellectual property office of the people's republic of China, First Search, Chinese Application 201310626100.0, June 17 2016, page 1

If you wish to add additional non-patent literature document citation information please click the Add button

Add

EXAMINER SIGNATURE

Examiner Signature

Date Considered

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

Application Number
6413

15061021

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Filing Date

2016-03-04

First Named Inventor

Matthew A. Markley

Art Unit

2453

Examiner Name

NGUYEN, TU T

Attorney Docket Number

YOR920120450US2

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J.Otterstedt/	Date (YYYY-MM-DD)	2017-03-14
Name/Print	Paul J.Otterstedt	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

#: 6415

(19) 中华人民共和国国家知识产权局



(12) 发明专利申请

(10) 申请公布号 CN 102387197 A

(43) 申请公布日 2012.03.21

(21) 申请号 201110308314.4

(22) 申请日 2011.09.29

(30) 优先权数据

12/895,685 2010.09.30 US

(71) 申请人 微软公司

地址 美国华盛顿州

(72) 发明人 M·洛利卡 S·萨姆帕斯

A · 帕勒卡 O · B · 伊万诺瓦

D·L·格林 I·本-沙哈尔

(74) 专利代理机构 上海专利商标事务所有限公

司 31100

代理人 杨洁

(51) Int. Cl.

H04L 29/08 (2006.01)

G06F 9/455(2006.01)

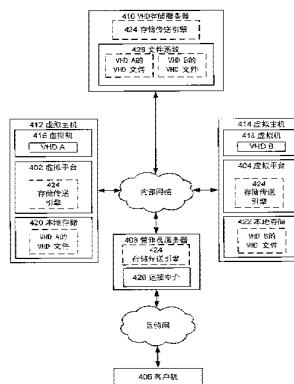
权利要求书 3 页 说明书 17 页 附图 14 页

(54) 发明名称

从服务器到主机的流虚拟机技术

(57) 摘要

本发明涉及从服务器到主机的流虚拟机技术。此处公开了用于将商用服务器配置成主机虚拟硬盘的技术。在示例性实施例中，虚拟硬盘文件可被分成多个差异 VHD 文件，当虚拟主机运行存储在服务器上的 VHD 文件时，文件中的一个或多个可被下载到虚拟主机。在一个或多个 VHD 文件被下载之后，虚拟主机可被配置为使用本地副本代替商用服务器上的副本。除以上之外，权利要求书、具体实施方式和附图中描述了其他技术。



1. 一种用于将虚拟硬盘驱动器部署到虚拟机主机的系统,包括:

第一服务器(412),包括:

第一处理器(102);以及

耦合至所述第一处理器(102)的第一存储器(106和/或104),所述第一存储器(106和/或104)包括在执行时使所述第一服务器执行以下动作的指令:

执行存储在第二服务器(410)上的虚拟硬盘文件集合(VHD A的VHD文件,(图5的504、506和508),(图6的505和508)或(图7的504和506))的客操作系统(220);

在同时执行虚拟硬盘文件集合(VHD A的VHD文件,(图5的504、506和508),(图6的505和508)或(图7的504和506))的所述客操作系统(220)时,存储附连于所述第一服务器(412)的存储设备(420)中的虚拟硬盘文件集合(VHD A的VHD文件,(图5的504、506和508),(图6的505和508)或(图7的504和506))的已接收的副本;以及

响应于确定所述虚拟硬盘文件集合(VHD A的VHD文件,(图5、6或7的510、512和514))的所述副本已被存储在附连于所述第一服务器(412)的所述存储设备(420)中,执行所述虚拟硬盘文件集合(VHD A的VHD文件,(图5、6或7的510、512和514))的所述副本的所述客操作系统(220)。

2. 如权利要求1所述的系统,其特征在于,所述第一存储器还包括在执行时使所述第一服务器执行以下动作的指令:

从所述虚拟硬盘文件集合引导所述客操作系统,其中所述虚拟硬盘文件集合包括主虚拟硬盘文件、机器差异虚拟硬盘文件和用户差异虚拟硬盘文件。

3. 如权利要求1所述的系统,其特征在于,所述第一存储器还包括在执行时使所述第一服务器执行以下动作的指令:

从所述虚拟硬盘文件集合以及存储在附连于所述第一服务器的所述存储设备中的主虚拟硬盘文件的预先高速缓存的副本引导所述客操作系统。

4. 如权利要求1所述的系统,其特征在于,所述第一存储器还包括在执行时使所述第一服务器执行以下动作的指令:

从所述虚拟硬盘文件集合以及存储在附连于所述第一服务器的所述存储设备中的用户差异虚拟硬盘文件引导所述客操作系统。

5. 如权利要求1所述的系统,其特征在于,所述第一存储器还包括在执行时使所述第一服务器执行以下动作的指令:

将用户差异虚拟硬盘文件合并到机器差异虚拟硬盘文件中,并将所述机器差异虚拟硬盘文件发送到所述第二服务器。

6. 如权利要求1所述的系统,其特征在于,还包括:

所述第二服务器,包括:

第二处理器;以及

耦合至所述第二处理器的第二存储器,所述第二存储器包括在执行时使所述第二服务器执行以下动作的指令:

将从所述第一服务器接收的用户差异虚拟硬盘文件合并到存储在所述第二服务器上的机器差异虚拟硬盘文件中。

7. 如权利要求1所述的系统,其特征在于,还包括:

所述第二服务器,包括:

第二处理器;以及

耦合至所述第二处理器的第二存储器,所述第二存储器包括在执行时使所述第二服务器执行以下动作的指令:

从所述第一服务器接收至少用户差异虚拟硬盘文件;以及

将至少所述用户差异虚拟硬盘文件发送到第三服务器。

8. 如权利要求 1 所述的系统,其特征在于,还包括:

所述第二服务器,包括:

第二处理器;以及

耦合至所述第二处理器的第二存储器,所述第二存储器包括在执行时使所述第二服务器执行以下动作的指令:

将所述虚拟硬盘文件集合的至少部分预先高速缓存到所述第二服务器的随机访问存储器中。

9. 如权利要求 1 所述的系统,其特征在于,还包括:

第四服务器,包括:

第四处理器;以及

耦合至所述第四服务器的第四存储器,所述第四存储器包括在执行时使所述第四服务器执行以下动作的指令:

响应于确定所述第一服务器包括机器差异虚拟硬盘文件的当前副本,指示客户机计算机系统连接到所述第一服务器。

10. 如权利要求 1 所述的系统,其特征在于,还包括:

第四服务器,包括:

第四处理器;以及

耦合至所述第四服务器的第四存储器,所述第四存储器包括在执行时使所述第四服务器执行以下动作的指令:

响应于确定所述第一服务器上的负载为低,指示客户机计算机系统连接到所述第一服务器。

11. 一种用于流虚拟硬盘的第一服务器 (412) 方法,包括:

从存储在第二服务器 (410) 中的机器差异虚拟硬盘文件 (图 6 的 506) 和存储在附连于所述第一服务器 (412) 的存储设备 (420) 中的主虚拟硬盘文件 (图 6 的 510) 引导客操作系统 (220),其中所述主虚拟硬盘文件 (510) 包括操作系统的副本,并且所述机器差异虚拟硬盘文件 (506) 包括用于为用户定制所述操作系统的所述副本的机器识别信息;

在同时引导所述客操作系统 (220) 时,从所述第二服务器 (410) 下载所述机器差异虚拟硬盘文件的副本 (图 6 的 512);以及

响应于确定所述机器差异虚拟硬盘文件的所述副本 (图 6 的 512) 已被下载,从所述机器差异虚拟硬盘文件的所述副本 (图 6 的 512) 和所述主虚拟硬盘文件 (图 6 的 510) 执行所述客操作系统 (220)。

12. 如权利要求 11 所述的方法,其特征在于,还包括:

将数据块写入被预先高速缓存到附连于所述第一服务器的所述存储设备中的用户差

异虚拟硬盘文件。

13. 如权利要求 11 所述的方法,其特征在于,还包括:

将存储在用户差异虚拟硬盘文件中的数据写入存储在第二服务器中的所述机器差异虚拟硬盘文件。

14. 如权利要求 11 所述的方法,其特征在于,还包括:

在迁移操作期间,将用户差异虚拟硬盘文件发送到所述第二服务器。

15. 如权利要求 11 所述的方法,其特征在于,从所述存储服务器下载所述机器差异虚拟硬盘文件的所述副本还包括:

使用空闲带宽从所述第二服务器下载所述机器差异虚拟硬盘文件的所述副本。

从服务器到主机的流虚拟机技术

技术领域

[0001] 本发明涉及虚拟机，尤其是从服务器到主机的流虚拟机。

背景技术

[0002] 虚拟机平台通过在每一操作系统自己的虚拟机内运行该操作系统来允许多个客操作系统在物理机上的同时执行。可以在虚拟机中提供的一个示例性服务是虚拟桌面会话。虚拟桌面会话在本质上是在虚拟机内运行的将其用户界面发送到远程计算机的个人计算机环境。为了虚拟化桌面，操作系统以及用户的程序和数据的全部被封装在虚拟硬盘文件（VHD）中，该虚拟硬盘文件随后由虚拟机中的虚拟处理器执行并由远程客户机设备访问。这一体系结构类似于远程桌面环境，但不是使多个用户同时连接到操作系统，而是每个用户能够访问在虚拟桌面环境里的虚拟机中执行的他们自己的操作系统。

[0003] 当在数据中心中部署虚拟桌面时，通常在诸如存储区域网络（SAN）目标的集中式存储服务器上部署虚拟硬盘文件，并且虚拟主机从 SAN 目标中读取虚拟硬盘文件。该部署允许虚拟桌面在数据中心的任意虚拟主机上被动态地实例化。SAN 目标使诸如磁盘的逻辑单元号（LUN）通过高速网络对虚拟主机可用，并且与网络附连存储不同，SAN 目标不作为对虚拟主机的远程存储出现。但是，SAN 目标是昂贵的，并且如果可以使用（带有低价的集中存储的）商用服务器来集中地部署虚拟硬盘文件将是有益的。

发明内容

[0004] 示例性实施例描述了用于将虚拟硬盘驱动器部署到虚拟机主机的系统。在所示实施例中，系统可包括但不限于第一服务器，该第一服务器被配置为执行存储在第二服务器上的虚拟硬盘文件集合的客操作系统；在执行虚拟硬盘文件集合的客操作系统的同时，将虚拟硬盘文件集合的已接收副本存储到附连于第一服务器的存储设备中；并且响应于确定虚拟硬盘文件集合的副本存储在附连于第一服务器的存储设备中，执行虚拟硬盘文件集合的客操作系统。除以上之外，权利要求书、具体实施方式和附图中描述了其他技术。

[0005] 除计算机系统以外，示例性实施例提供了将虚拟硬盘驱动器部署到虚拟机主机的操作过程。示例性操作过程包括但不限于，从存储在第二服务器中的机器差异虚拟硬盘文件和存储在附连于第一服务器的存储设备中的主虚拟硬盘文件引导客操作系统，其中主虚拟硬盘文件包括操作系统的副本，并且机器差异虚拟硬盘文件包括用于为用户定制操作系统的副本的机器识别信息；在同时引导客操作系统时，从第二服务器下载机器差异虚拟硬盘文件的副本；并且，响应于确定机器差异虚拟硬盘文件的副本已被下载，从机器差异虚拟硬盘文件的副本和主虚拟硬盘文件执行客操作系统。除以上之外，权利要求书、具体实施方式和附图中描述了其他技术。

[0006] 在另一个示例中，包括可用于将虚拟硬盘驱动器部署到虚拟机主机的可执行指令的计算机可读存储介质。示例性计算机可读存储介质可包括在执行时使虚拟主机执行以下动作的指令：将客操作系统从虚拟硬盘文件集合引导到虚拟机，其中集合的第一部分被存

储到附连于第一服务器的存储设备上,第二部分被存储到第二服务器上;在引导客操作系统时,下载集合第二部分的至少子部分;并且响应于确定集合第二部分的至少子部分被下载到附连于第一服务器的存储设备,执行附连于第一服务器的存储设备上的客操作系统。除以上之外,权利要求书、具体实施方式和附图中描述了其他技术。

[0007] 本领域技术人员将理解,此处所描述的一个或多个方面可包括但不限于用于实现此处所引用的方面的线路和/或编程;该线路和/或编程实质上可以是配置成实现本文所述方面的硬件、软件和/或固件的任何组合,这取决于系统设计者的设计选择。

[0008] 以上是概述,并且因此必然包含细节的简化、一般化及省略。本领域技术人员将明白,本概述只是说明性的并且决不旨在是限制性的。

附图说明

- [0009] 图 1 描述了示例计算机系统。
- [0010] 图 2 描绘了描述示例性虚拟平台的操作环境。
- [0011] 图 3 描绘了描述示例性虚拟平台的操作环境。
- [0012] 图 4 描绘了数据中心的高级框图。
- [0013] 图 5 描绘了虚拟主机和 VHD 存储服务器的高级框图。
- [0014] 图 6 描绘了虚拟主机和 VHD 存储服务器的高级框图。
- [0015] 图 7 描绘了虚拟主机和 VHD 存储服务器的高级框图。
- [0016] 图 8 描绘了虚拟主机和 VHD 存储服务器的高级框图。
- [0017] 图 9 描绘了两个虚拟主机和 VHD 存储服务器的高级框图。
- [0018] 图 10 描绘了操作过程。
- [0019] 图 11 描绘了图 10 的操作过程的替换实施例。
- [0020] 图 12 描绘了操作过程。
- [0021] 图 13 描绘了图 12 的操作过程的替换实施例。
- [0022] 图 14 描绘了操作过程。
- [0023] 图 15 描绘了图 14 的操作过程的替换实施例。

具体实施方式

[0024] 本发明可使用一个或多个计算机系统。图 1 及以下讨论旨在提供可在其中实现所公开的主题的合适计算环境的简要概括描述。

[0025] 贯穿本说明书使用的术语线路可包括诸如硬件中断控制器、硬盘驱动器、网络适配器、图形处理器、基于硬件的视频/音频编解码器等硬件组件,以及用于操作这些硬件的固件。术语线路还可包括由固件和/或软件来配置的微处理器、应用程序专用集成线路、以及处理器,例如多核通用处理单元的核(被配置为读取并执行指令的单元)。处理器可由从例如 RAM、ROM、固件和/或大容量存储等存储器加载的指令来配置,从而具体化可用于配置处理器来执行功能的逻辑。在其中线路包括硬件和软件的组合的示例实施例中,实现者可以编写具体化逻辑的源代码,该源代码随后被编译成可由硬件执行的机器可读代码。因为本领域技术人员可以明白现有技术已经进化到硬件实现的功能或软件实现的功能之间几乎没有差别的地步,因而选择硬件还是软件来实现此处描述的功能只是一个设计选择。换

言之,由于本领域的技术人员可以理解软件进程可被变换成等效的硬件结构,且硬件结构本身可被变换成等效的软件进程,因此选择硬件实现或是软件实现留给了实现者。

[0026] 现在参考图 1,描绘了示例性计算系统 100。计算机系统 100 可包括处理器 102,例如,执行核。尽管示出了一个处理器 102,但在其他实施例中,计算机系统 100 可具有多个逻辑处理器,例如每一处理器衬底有多个执行核,和 / 或各自可具有多个执行核的多个处理器衬底。如图 1 所示,各种计算机可读存储介质 110 可由一个或多个系统总线互连,系统总线将各种系统组件耦合到处理器 102。系统总线可以是几种类型的总线结构中的任何一种,包括存储器总线或存储器控制器、外围总线、以及使用各种总线体系结构中的任一种的局部总线。在示例实施例中,计算机可读存储介质 110 可以包括例如随机存取存储器 (RAM) 104、存储设备 106 (例如电机硬盘驱动器、固态硬盘驱动器等)、固件 108 (例如闪速 RAM 或 ROM)、以及可移动存储设备 118 (例如 CD-ROM、软盘、DVD、闪速驱动器、外部存储设备等)。本领域的技术人员应当理解,可使用其他类型的计算机可读存储介质,如磁带盒、闪存卡、和 / 或数字视频盘。

[0027] 计算机可读存储介质 110 可提供对计算机 100 的处理器可执行指令 122、数据结构、程序模块和其他数据的非易失性和易失性存储。基本输入 / 输出系统 (BIOS) 120 可被存储在固件 108 中,它包含帮助在诸如启动期间在计算机系统 100 内的各元件之间拷贝信息的基本例程。包括操作系统和 / 或应用程序的多个程序可被存储在固件 108、存储设备 106、RAM 104 和 / 或可移动存储设备 118 上,并且可由处理器 102 来执行。

[0028] 命令和信息可由计算机 100 通过输入设备 116 来接收,输入设备可包括但不限于键盘和指向设备。其他输入设备可以包括话筒、操纵杆、游戏垫、扫描仪等等。这些和其他输入设备常通过耦合到系统总线的串行端口接口连接到处理器 102,但也可通过其他接口连接,如并行端口、游戏端口或通用串行总线 (USB)。显示器或其他类型的显示设备也可经由诸如视频适配器等可以是图形处理单元 112 的一部分或可连接到图形处理单元 112 的接口来连接到系统总线。除显示器之外,计算机通常包括其他外围输出设备,如扬声器和打印机 (未示出)。图 1 的示例性系统还可包括主适配器、小型计算机系统接口 (SCSI) 总线和连接到 SCSI 总线的外部存储设备。

[0029] 计算机系统 100 可使用至一个或多个远程计算机,如远程计算机的逻辑连接在网络化环境中操作。远程计算机可以是另一计算机、服务器、路由器、网络 PC、对等设备或其他常见的网络节点,并且通常包括上面关于计算机系统 100 所述的许多或全部元件。

[0030] 当在 LAN 或 WAN 联网环境中使用时,计算机系统 100 可通过网络接口卡 114 连接至 LAN 或 WAN。NIC 114 可以是内置或外置的,其可以连接到系统总线。在联网环境中,相对于计算机系统 100 所描述的程序模块或其部分可被储存在远程存储器存储设备中。可以理解,所描述的网络连接是示例性的,且可以使用在计算机之间建立通信链路的其他手段。虽然可构想本发明的许多实施例尤其适用于计算机化的系统,然而在本说明中不意味着将本发明限于那些实施例。

[0031] 转向图 2,示出了可被用来生成虚拟机的示例性虚拟平台。在该实施例中,系统管理程序微内核 202 可被配置成控制并仲裁对计算机系统 200 的硬件的访问。系统管理程序微内核 202 可以隔离一个分区中的进程,使其不能访问另一分区的资源。例如,系统管理程序微内核 202 可以生成称为分区的执行环境,如子分区 1 到子分区 N (其中 N 是大于 1 的整

数)。在该实施例中,子分区是系统管理程序微内核 202 支持的基本隔离单元。每一子分区可被映射到在系统管理程序微内核 202 控制之下的一组硬件资源,例如,存储器、设备、处理器周期等。在各实施例中,系统管理程序微内核 202 可以是独立的软件产品、操作系统的一部分、嵌入在主板的固件中、专用集成线路、或其组合。

[0032] 系统管理程序微内核 202 可以通过限制物理计算机系统上的客操作系统的存储器视图来实施分区划分。当系统管理程序微内核 202 实例化一虚拟机时,它可以将系统物理存储器 (SPM) 的页 (例如,具有开始和结束地址的固定长度存储器块) 分配给虚拟机作为客物理存储器 (GPM)。在该实施例中,客机的受限的系统存储器视图受系统管理程序微内核 202 控制。术语客物理存储器是从虚拟机的视点描述存储器页的简写方式,且术语系统物理存储器是从物理系统的观点描述存储器页的简写方式。因此,被分配给虚拟机的存储器页会有客物理地址 (虚拟机所使用的地址) 和系统物理地址 (页的实际地址)。

[0033] 客操作系统可以虚拟化客物理存储器。虚拟存储器是一种管理技术,其允许操作系统过度提交存储器,并且给予应用程序对连续工作存储器的唯一访问。在虚拟化环境中,客操作系统可以使用一个或多个页表来将已知为虚拟客地址的虚拟地址转换成客物理地址。在该示例中,存储器地址可以具有客虚拟地址、客物理地址以及系统物理地址。

[0034] 在所描绘的示例中,父分区组件也可被认为是类似于 Xen 的开源系统管理程序的域 0,父分区组件可以包括主机 204。主机 204 可以是操作系统 (或一组配置实用程序),并且主机 204 可以被配置成通过使用虚拟服务提供者 228 (VSP) 向在子分区 1-N 中执行的客操作系统提供资源。VPS 228 (一般在开源社区中被称为后端驱动程序) 可用来通过虚拟服务客户机 (VSC) (在开源社区或类虚拟化设备中一般称为前端驱动程序) 对到硬件资源的接口进行多路复用。如图所示,虚拟服务客户机可以在客操作系统的上下文中执行。然而,这些驱动程序不同于客机中的其余驱动程序,因为它们向它们提供了系统管理程序而非客机。在一示例性实施例中,虚拟服务提供者 228 与虚拟服务客户机 216 和 218 通信所使用的路径可以被视为虚拟路径。

[0035] 如图所示,仿真器 234 (例如虚拟化 IDE 设备、虚拟化视频适配器、虚拟化 NIC 等) 可被配置成在主机 204 中运行并被附连到可供客操作系统 220 和 222 使用的资源。例如,当客操作系统接触 (被映射到设备的寄存器所处的) 存储器位置或被映射到设备的存储器时,微内核系统管理程序 202 可截取该请求并将客机试图的值传递给相关联的仿真器。在该示例中,资源可被认为是虚拟设备所处的位置。该例中仿真器的使用可以被视为仿真路径。仿真路径与虚拟路径相比是低效的,因为与在 VSP 和 VSC 之间传递消息相比,它需要更多的 CPU 资源来仿真设备。例如,可以将经由仿真路径把值写入盘所需的、被映射至寄存器的存储器上的几百个动作减少为一个消息,该消息在虚拟路径中从 VSC 被传递至 VSP。

[0036] 每一子分区可包括一个或多个虚拟处理器 (230 和 232),客操作系统 (220 和 222) 可管理并调度线程以便在这些虚拟处理器上执行。一般而言,虚拟处理器是提供具有特定架构的物理处理器的表示的可执行指令以及相关联的状态信息。例如,一个虚拟机可具有带有英特尔 x86 处理器特性的虚拟处理器,而另一虚拟处理器可具有 PowerPC 处理器的特性。本例中的虚拟处理器可被映射到计算机系统的逻辑处理器,使得实现虚拟处理器的指令将受到逻辑处理器的支持。由此,在包括多个逻辑处理器的实施例中,在例如其他逻辑处理器执行系统管理程序指令时,虚拟处理器可以由逻辑处理器同时执行。分区中虚拟处理

器和存储器的组合可被认为是虚拟机。

[0037] 客操作系统 (220 和 222) 可以是任何操作系统,如来自**微软®**、**苹果®**、开源社区等的操作系统。客操作系统可包括用户 / 内核操作模式,并且可具有能包括调度器、存储器管理器等的内核。一般而言,内核模式可包括处理器中的执行模式,该执行模式授予至少对特权处理器指令的访问。每一客操作系统可具有相关联的文件系统,该文件系统上存储有诸如终端服务器、电子商务服务器、电子邮件服务器等应用程序以及客操作系统本身。客操作系统可调度线程来在虚拟处理器上执行,并且可实现此类应用程序的实例。

[0038] 现在参考图 3,示出了上述图 2 中的替换虚拟平台。图 3 描绘了与图 2 的组件相类似的组件;然而,在该示例实施例中,系统管理程序 302 可包括微内核组件以及和图 2 的主机 204 中的组件(如虚拟服务提供者 228 和设备驱动程序 224)相类似的组件,而管理操作系统 304 可包含例如用于配置系统管理程序 302 的配置实用程序。在该体系结构中,系统管理程序 302 可以执行与图 2 的系统管理程序微内核 202 相同或相似的功能;然而,在该体系结构中,系统管理程序 302 可被配置成向在子分区中执行的客操作系统提供资源。图 3 的系统管理程序 302 可以是独立的软件产品、操作系统的一部分、嵌入在主板的固件内,或者系统管理程序 302 的一部分可以由专用集成线路来实现。

[0039] 现在转向图 4,它示出了示例性数据中心的高级框图。图 4 一般地示出了将虚拟桌面部署到客户机的环境。简言之,图 4 中示出的某些元件是用虚线指示的,其指示了认为它们在其所示位置内可选地放置和 / 或可以以分布式的方式来实现它们。因此,存储传送引擎 424 可位于管理员服务器 408 和 / 或 VHD 存储服务器 410 上。以下描述将调用如何以及何处实现存储传送引擎 424 的具体示例,然而,本发明不限于这些具体的样例。

[0040] 在所示示例中,数据中心可以包括内部网络,其将多个虚拟主机 (412 和 414) 耦合到管理员服务器 408 和虚拟硬盘驱动器 (VHD) 存储服务器 410,虚拟主机可包括类似于图 2 或 3 所示组件的组件。本领域的技术人员可以理解,尽管示出了两个虚拟主机,但数据中心可具有更多个虚拟主机。同时,尽管每个虚拟主机 (412 和 414) 被示为运行一个虚拟机 (416 和 418),但是每个虚拟主机可以同时主存更多的虚拟机。或者换言之,数据中心可以有 M 个虚拟主机(其中 M 是大于 1 的整数),并且 M 个虚拟主机中的每一个可以主存 N 个虚拟机(其中 N 也是大于 1 的整数)。虚拟平台 402 和 404 是如上在图 2 和图 3 中所述的虚拟基础结构组件的逻辑抽象。在以下章节中被描述为在虚拟平台 402 或 404 “以内”的功能可以在图 2 和图 3 所述的一个或多个元件中实现。

[0041] 可包括类似于图 1 的计算机系统 100 的组件的管理员服务器 408 还可包括用于连接中介 426 的线路,例如由从计算机可读存储介质载入的指令来配置的处理器。连接中介 426 可从客户机接收虚拟桌面请求,并可选择主机来运行用户的虚拟机。管理员服务器 408 可将用户的配置文件发送到虚拟主机 412。虚拟主机 412 可包括在虚拟平台 402 内实现的虚拟机引擎(未示出),该虚拟机引擎可根据配置文件中的信息执行并实例化虚拟机 416。作为实例化过程的一部分及如以下段落中详细描述,虚拟平台 402 可被配置为,为虚拟机 416 确定虚拟硬盘文件的位置,以便引导存储其上的客操作系统。

[0042] VHD 存储服务器 410 可包括类似于图 1 的计算机系统 100 的组件。在示例性实施例中,可从诸如通用处理单元、1 或 100GB 网络适配器、串行高级技术附件 (SATA) 硬盘、小型计算机系统接口 (SCSI) 硬盘的商用部件来构建 VHD 存储服务器 410,而不是在昂贵的 SAN 上

实现。VHD 存储服务器 410 可载入可控制计算机系统的硬件的例如 Windows Server 2008® 的商用操作系统。在此示例中,操作系统可执行用于实现应用层网络协议的指令,应用层网络协议提供对存储在由操作系统控制的文件系统 428 中的文件的共享访问。VHD 存储服务器 410 在内部网络上共享其文件系统 428,而不像 SAN 那样展示 LUN。这样,在该示例性实施例中,虚拟主机可将 VHD 存储服务器 410 作为联网存储设备来检测。操作中,虚拟主机发送读取 VHD 存储服务器的文件系统 428 中的文件的请求,并且 VHD 存储服务器 410 上运行的操作系统检索所期望的数据。

[0043] VHD 存储服务器 410 可在其文件系统 428 中存储多个虚拟硬盘文件。简言之,虚拟硬盘文件是物理硬盘的虚拟表示并且存储数据块,例如,表示诸如客操作系统的数据结构和存储在客机的文件系统(文件、文件夹等)中的数据的数据的字节或位序列。虚拟平台可接收已从客机发出的输入/输出(I/O)任务,并代表客机来执行它们。但是,从 VHD 存储服务器 410 引导和运行客 OS 可能导致严重的等待时间问题,因为未针对该部署优化 VHD 存储服务器 410。例如,当虚拟平台 402 必须读取 VHD 存储服务器 410 的 VHD 字节时,引导客机可能需要超过一分钟,例如从引导加载器开始到其退出的时间所测量的。如果多个用户试图同时登录到他们的虚拟桌面时,该操作将需要明显更长的时间。为了加速该过程,在示例性实施例中,每个 VHD 可被分解为多个差异 VHD 文件,并且存储传送引擎 424 可将不同的文件(或其部分)选择性地高速缓存到数据中心的各位置中。

[0044] 在示例性实施例中,表示虚拟硬盘的数据可被分解为多个文件,并且数据的组合构成包括客操作系统的虚拟硬盘。差异虚拟硬盘文件将虚拟硬盘的当前状态表示为与父图像相比一组已修改的块。父图像可以是另一个差异虚拟硬盘文件或者基础虚拟硬盘文件。由于这一点,每个差异虚拟硬盘文件可包括唯一的标识符、其父的标识符和标识其父的位置的值。每个差异 VHD 文件可被链接在一起,并且链接链可指向基础虚拟硬盘文件的位置。

[0045] 在示例性实施例中,每个虚拟硬盘可被分解为三个虚拟硬盘文件。尽管示出了三个虚拟硬盘文件,但本发明并不限于此配置,并且虚拟硬盘可被分解为两个虚拟硬盘文件或更多个。在示例性实施例中,可从主 VHD 文件、机器差异 VHD 文件和用户差异 VHD 文件中生成虚拟硬盘文件。主 VHD 文件可包括类似于在 OS 安装 DVD 上找到且出售给消费者的操作系统的副本。主 VHD 文件还可包括数据中心的由用户使用的程序。例如,如果正向需要这样的软件的用户提供数据中心,则主 VHD 文件可与办公室生产性软件在一起。在另一个示例中,如果所预期的用户是 web 设计者,则主 VHD 文件可包括 web 设计软件。

[0046] 机器差异 VHD 文件可被链接到主 VHD 文件。实际上,机器差异 VHD 文件可包括向特定用户提供主 VHD 文件的数据块。例如,机器差异 VHD 文件可包括将一个用户的计算环境与另一个相区分的数据,诸如机器标识符、活动目录标识符等。在示例性实施例中,数据中心中可存在用于每个唯一的用户账户的机器差异 VHD 文件。在相同或另一个实施例中,每个主 VHD 文件可被链接到相同的主 VHD 文件。

[0047] 在另一个示例性实施例中,可以可选地使用用户差异 VHD 文件。在使用用户差异 VHD 文件的示例实施例中,它可被链接到主 VHD 文件并用于存储在会话期间被写到磁盘的数据块。例如,当用户与虚拟桌面交互时,客操作系统将信息写到磁盘。在示例性实施例中,这些数据块可被存储到用户差异 VHD 文件中。在相同或另一个实施例中,在快照期间写入磁盘的数据可被存储到用户差异 VHD 文件中。这样,用户差异 VHD 文件可以是可写的,而机

器差异 VHD 文件和主 VHD 文件被设为只读。在该示例性配置中,当虚拟机运行时仅有用户差异 VHD 文件中的数据改变。

[0048] 在示例性配置中,链中最后的虚拟硬盘文件的位置可被存储到虚拟机的配置文件中。当虚拟平台接收请求以实例化虚拟机时,虚拟平台可读取配置文件以发现链中最后链接的位置。虚拟平台可随后发现其父的位置并打开父。虚拟平台可继续发现链中的虚拟硬盘文件直到它发现基础虚拟硬盘文件的位置并将其打开。

[0049] 在另一个示例性实施例中,链中最后的只读虚拟硬盘文件的位置可被存储到配置文件中,例如,机器差异 VHD 文件。在该示例中,当引导虚拟机时可动态地创建用户差异 VHD 文件。例如,响应于检测到指示只读文件的父位置值,例如虚拟机引擎的虚拟平台 402 中的处理可创建用户差异 VHD 文件。在该示例中,虚拟机引擎可创建用户差异 VHD 文件并将其动态地链接到只读文件。

[0050] 虚拟平台可使用位图和预先配置的操作次序来读取并写入虚拟硬盘文件集合。例如,每个虚拟硬盘文件可包括位图和数据。位图指示哪些数据块位于差异磁盘 (1) 中而哪些数据块没有 (0)。在该示例性实施例中,当虚拟平台 (即,虚拟平台中运行的存储服务) 试图从文件集合进行读取时,它可被配置为首先试图从用户差异虚拟硬盘文件中读取,并且如果块不存在,则它可拷贝到机器差异虚拟硬盘文件。在任一文件中都没有该块的情况下,则虚拟平台 402 可被配置为从主虚拟硬盘文件中读取该块。对于写操作,虚拟平台 402 可被配置为写入链中的第一个链接 (例如,用户差异 VHD 文件) 并将文件位图中的对应块标记为 1。对该数据块的此后的读 IO 任务可被定向到用户差异 VHD 文件。

[0051] 转回到图 4,在示例性实施例中,存储传送引擎 424 可被配置为当虚拟机在运行时选择性地 VHD 文件下载到虚拟主机本地存储。在虚拟硬盘文件被下载后,例如图 2 的微内核系统管理程序 202 的虚拟平台可使用“Setparent (设定父)”命令来重新配置每个 VHD 文件和 / 或配置文件中的父位置值以指向文件的本地副本。通过选择性地下 VHD 文件,发生在 VHD 存储服务器 410 上的磁盘操作的数量被最小化并且处理用于 M 个虚拟机的存储操作的压力被消除。

[0052] 在用户退出之后,对用户差异 VHD 文件的写可被丢弃 (在汇合部署中) 或与例如机器差异 VHD 文件进行合并 (在个人虚拟机部署中)。合并操作可在数据中心的多个位置之一进行。在第一个实施例中,可在虚拟主机上合并两个文件的本地副本。在第二个实施例中,可在 VHD 存储服务器 410 上合并两个文件的本地副本。在第三个实施例中,运行在虚拟主机上的存储传送引擎 424 将本地存储 420 中的用户差异 VHD 文件的副本与存储在 VHD 存储服务器 410 中的机器差异 VHD 文件相合并。在第四个实施例中,运行在 VHD 存储服务器 410 上的存储传送引擎 424 可将本地存储 420 中的用户差异 VHD 文件的副本与存储在 VHD 存储服务器 410 中的机器差异 VHD 文件相合并。合并操作的位置可由服务器管理员来选择。

[0053] 无论拷贝操作在何处进行,存储传送引擎 424 可通过将用户差异 VHD 文件中的数据块拷贝到机器差异 VHD 文件的对应块中来合并存储在用户差异 VHD 文件中的差异与机器差异 VHD 文件。VHD 存储服务器 410 可接收已修改的机器差异 VHD 文件的副本,并使用已修改的版本来重写文件的副本。在此处,例如虚拟主机 412 的虚拟主机可删除其副本或在例如本地存储 420 的本地存储中继续对其进行存储。在将其存储在本地存储 420 的实例中,

可向管理员服务器 408 发送指示机器差异 VHD 文件的副本的时戳的信号。在用户再次登陆的情况下,连接中介 426 可将存储在本地存储 420 中的机器差异 VHD 文件的时戳与 VHD 存储 410 中对应文件的时戳进行比较。在时戳相匹配的情况下,连接中介 426 可被配置为试图在虚拟主机 412 上而非数据中心的不同的主机上实例化虚拟机 416。

[0054] 图 5-8 示出不同的部署配置。接着要实现对部署配置的选择,但是,可通过考虑内部网络的可用带宽以及将 VHD 存储服务 410 的处理器连接到物理存储设备(例如,内部 SATA 总线带宽)的内部总线的可用带宽来进行选择。例如,网络管理员可测试他或她的数据中心以观察从存储有全部 VHD 文件的 VHD 存储服务器 410 引导操作系统所需的时间。如果(基于偏好)所用时间过长,则管理员可使用此处描述的不同的部署配置中的一个。在管理员希望最大化性能的实例中,他或她可使用最小化必须通过内部网络发送的数据量的部署。

[0055] 在相同或另一个实施例中,存储传送引擎 424 可被配置为基于算法选择性地拷贝某些文件,该算法最小化用于拷贝 VHD 文件的可用带宽量。在相同或另一个实施例中,算法可被配置为检测哪个文件处理最多的 IO 任务并拷贝该文件。

[0056] 转向图 5,它示出了示例性部署。在所示实施例中,在由图 4 的连接中介 426 作出在虚拟主机 412 上实例化虚拟机 416 的决定之后,存储传送引擎 424 已被配置为将主 VHD 文件 504、机器差异 VHD 文件 506 和用户差异 VHD 文件 508 拷贝到本地存储 420。此图及以下各图中的虚线示出主 VHD 文件 504、机器差异 VHD 文件 506 和用户差异 VHD 文件 508 的副本可从 VHD 存储服务器 410 拷贝到虚拟主机 412。

[0057] 在示例性实施例中,拷贝操作可在引导之前开始。例如,虚拟平台 402 一确定文件集合的位置操作就可开始。在另一个实施例中,当 VHD 存储服务器 410 上的存储传送引擎 424 接收请求以从虚拟平台 402 打开文件集合时,拷贝操作可开始。

[0058] 在两种情形中,虚拟平台 402 可响应于接收配置文件来开始该过程。虚拟平台 402 可打开文件并读取可包括例如用户差异 VHD 文件 508 的网络路径的 VHD 位置字段,可以是 \\VHDstorage server\share\...\userXXXuserdiff.vhd。虚拟平台 402 可向 VHD 存储服务器 410 发送打开该文件的请求,并且 VHD 存储服务器 410 可处理该请求。虚拟平台 402 随后可读取指示用户差异 VHD 文件的父的位置的值,可以是 \\VHDstorage server\share\...\userXXXmachinediff.vhd。类似地,虚拟平台 402 随后可发现机器差异 VHD 文件的父位于 \\VHDstorage server\share\...\master.vhd。在图中这些链接中的每一个被示为箭头。

[0059] 在替换的实施例中,配置文件可包括机器差异 VHD 文件 506 的 VHD 位置值,即 \\VHDstorage server\share\...\userXXXmachinediff.vhd。在该示例中,虚拟平台 402 可确定机器差异 VHD 文件 506 是只读的并且生成用户差异 VHD 文件 508。虚拟平台 402 可将新生成的用户差异 VHD 文件的父位置设为指向机器差异 VHD 文件 506,并修改配置文件以指向新生成的用户差异 VHD 文件的位置。

[0060] 在文件被打开后,虚拟平台可读取位图并启用虚拟机 416 中的虚拟处理器。虚拟机 416 中的 OS 引导加载器程序可发出对于第一块的 IO 任务,需要读取该第一块以便引导客操作系统。虚拟平台 402 可使用位图来确定哪个磁盘包括所请求的字节,从 VHD 存储服务器 410 上打开的文件取回字节,并将字节返回给引导加载器。

[0061] 在 VHD 链被存储到本地存储 420 之后,可以向虚拟平台 402 发送信号以重置存

储在 VHD 文件中的链接和配置文件。当虚拟机 416 运行时,虚拟平台 402 可修改配置文件中的父位置以指向已下载的副本,例如本地用户差异 VHD 文件 514,并且存储在机器差异 VHD 文件 512 的本地副本中的父位置信息可被重置为指向本地主 VHD 文件 510。例如,如微内核系统管理程序 202 的虚拟平台 402 可包括存储传送引擎 424 可以访问以便修改父位置信息的接口。在特定的示例中,可以访问接口并且配置文件中的父位置信息可从 \\VHDstorage server\share\...\userXXXuserdiff.vhd 变为 C:\ProgramData\...\Virtual Machines\userXXXuserdiff.vhd。在此处,存储传送引擎 424 可关闭到 VHD 存储服务器 410 的连接,并且将之后的 IO 任务发送到 VHD 文件的本地副本而不是存储在 VHD 存储服务器 410 上的副本。例如,当由虚拟平台 402 接收读 IO 任务时,已更新的位置信息可被读取并使虚拟平台 402 发送任务到本地存储 420

[0062] 由于链中的 VHD 文件可能非常大,例如主 VHD 文件 504 单独可有 7 或 8GB,在示例性实施例中,可使用快速文件拷贝技术来减少拷贝时间。例如,可以在虚拟主机 412 和 / 或 VHD 存储服务器 410 上执行后台智能传送服务 (BITS)。BITS 可使用空闲的网络带宽来拷贝文件并在后台拷贝数据。在操作中,BITS 程序可以监控 VHD 存储服务器 410 上的网络通信并扼制其自己的拷贝以确保其他下载获得所需的带宽。

[0063] 在相同或另一个实施例中,VHD 存储服务器 410 可被配置为将特定的字节块预先高速缓存到 RAM 中从而加速拷贝操作。例如,引导加载器通常需要来自主 VHD 文件 504 的相同的数据块以执行登录操作,在示例性实施例中,这些数据块可在 VHD 存储服务器 410 引导时被高速缓存到 RAM 中,以便 VHD 存储服务器 410 不必首先从大容量存储中取回字节。在另一个示例性实施例中,在 VHD 存储服务器 410 具有足够 RAM 的实例中,可以预先高速缓存整个主 VHD 文件 504。

[0064] 在示例性实施例中,主 VHD 文件 504 和机器差异 VHD 文件 506 可被标记为只读。这样,这些文件可以被拷贝到虚拟主机 406 而没有问题。但是,用户差异 VHD 文件 508 是读 / 写并且需要以确保用户数据不会丢失的方式来拷贝。在示例性实施例中,用户差异 VHD 文件 508 起始于 0 字节 (引导时文件为空) 并且大小快速增长。在该示例中,当存储传送引擎 424 启动拷贝操作时,可创建本地用户差异 VHD 文件 514,并且虚拟平台 402 可被配置为向用户差异 VHD 文件 508 和本地用户差异 VHD 文件 514 两者发送写 IO 任务。存储传送引擎 424 可存储指示块拷贝行的值,并且可以串行地开始向本地用户差异 VHD 文件 514 拷贝块,从用户差异 VHD 文件 508 的起始到结尾,而客操作系统向两个副本写入。在向本地存储 510 拷贝一块后,该值被递增。当该值到达用户差异 VHD 文件 508 的结尾时,两个文件已被同步,并且存储传送引擎 424 可以停止镜像对用户差异 VHD 文件 508 的写入。

[0065] 转向图 6,它示出主 VHD 文件的副本被预先高速缓存到虚拟主机 412 的替换实现。在此部署中,数据中心的每个虚拟主机可以使本地主 VHD 文件 510 被预先高速缓存,从而不必从 VHD 存储服务器 410 拷贝该文件。在此部署中,VHD 存储服务器 410 可以具有或不具有主 VHD 文件的副本,但是,它可包括副本使得被加至数据中心的虚拟主机可以下载副本。在此示例性配置中,机器差异 VHD 文件 506 可包括指向主 VHD 文件的本地副本的位置信息,从而消除了必须通过内部网络发送对主 VHD 文件的读操作的情况。与上面类似的,在 VHD 链被存储到本地存储 420 之后,可以向虚拟平台 402 发送信号以调整存储在 VHD 文件中的父位置信息和配置文件。

[0066] 转向图 7, 它示出了另一个替换的部署。在此示例性实施例中, 配置文件可指向机器差异 VHD 文件 506 和主 VHD 文件 504。在另一个配置中, 配置文件可指向本地用户差异 VHD 文件 514。在引导期间, 下载主 VHD 文件 504 的副本和机器差异 VHD 文件 506 时, 写可以被立即发送给本地用户差异 VHD 文件 514。与上面类似的, 在 VHD 链被存储到本地存储 420 之后, 可以向虚拟平台 402 发送信号以调整存储在 VHD 文件中的父位置信息和配置文件。

[0067] 图 8 示出了另一个替换的部署。根据该示例性实施例, 本地存储 420 可包括本地用户差异 VHD 文件 514 和本地主 VHD 文件 510, 并且 VHD 存储服务器 410 可包括机器差异 VHD 文件 506。在此示例中, 虚拟机 416 的配置文件可包括用户差异 VHD 文件 514 的本地副本的位置。本地用户差异 VHD 文件 514 可指向 VHD 存储服务器 410 上的机器差异 VHD 文件 506, 该机器差异 VHD 文件可指向本地主 VHD 文件 510。与上述操作类似的, 当虚拟机 416 正在运行时, 可将机器差异 VHD 文件 506 的副本下载到虚拟主机 406 并将其存储为本地机器差异 VHD 文件 512。与上面类似的, 在本地机器差异 VHD 文件 512 被存储到本地存储 420 之后, 可向虚拟平台 402 发送信号以调整存储在本地用户差异 VHD 文件中的链接, 以指向本地机器差异 VHD 文件 512 的文件系统位置。

[0068] 现在转向图 9, 它示出将虚拟机 416 从虚拟主机 412 迁移到虚拟主机 414 的操作环境。在所示实施例中, 假设当拷贝操作开始时, 虚拟硬盘 A 的 VHD 文件链已被下载到虚拟主机 406。但是, 情况可能并非如此, 并且当启动实况迁移操作时, 虚拟主机 412 可能具有 VHD 文件的任意组合。当启动存储传送引擎 424 时, 例如, 虚拟主机 412 可以开始拷贝确保将 VHD A 拷贝到虚拟主机 414 所必需的 VHD 文件。在特定的示例中, 这可以仅包括将本地用户差异 VHD 文件 514 拷贝到 VHD 存储服务器 410, 并且可以使用此处描述的镜像技术来拷贝文件。在需要被拷贝的文件被上载到 VHD 存储服务器 410 之后, 对 VHD A 的控制可从虚拟主机 412 切换到虚拟主机 414, 并且虚拟平台 404 可以开始将其需要的文件下载到本地存储 422。例如, 虚拟主机 414 可能已经包括本地主 VHD 文件 802 和 / 或本地机器差异 VHD 文件 804。这样, 在此示例中, 用户差异 VHD 文件 508 可被拷贝到虚拟主机 414 并被存储为本地用户差异 VHD 文件 806。类似地, 虚拟主机 404 可被配置为在下载结束后重置任意链接, 从而将读和写定向到 VHD 文件的本地副本而非 VHD 存储服务器 410 上的副本。

[0069] 以下是描绘操作过程的一系列流程图。为便于理解, 流程图被组织成使得初始流程图通过总体“大图”视点来演示各实现, 并且后续的流程图中提供以虚线示出的进一步的附加物和 / 或细节。此外, 本领域的技术人员可以理解, 虚线所描绘的操作过程被认为是任选的。

[0070] 现在转向图 10, 它示出始于操作 1000 的操作过程, 并且操作 1002 示出包括用于执行存储在第二服务器上的虚拟硬盘文件集合的客操作系统的线路的虚拟主机 412。例如, 虚拟平台 402 之前可能已经打开了集合, 例如存储在第二服务器上的一个或多个 VHD 文件, 例如存储在 VHD 存储服务器 410 的文件系统 428 中。虚拟平台 402 (例如可执行指令) 可以调度指示处理器上的虚拟处理器的线程, 并且处理器可以执行客操作系统, 该客操作系统向存储在文件系统 428 中的 VHD 文件集合发出 IO 任务。示例 IO 任务可以是读。在此示例中, 虚拟平台 402 可接收 IO 任务并且通过内部网络从 VHD 文件读取所请求的数据块并将字节发送回到客 OS。

[0071] 继续对图 10 的描述, 操作 1004 示出在实施例中, 虚拟主机可包括用于当同时执行

虚拟硬盘文件集合的客操作系统时,存储附连于第一服务器的存储设备中的虚拟硬盘文件集合的已接收的副本的线路。例如,转回到图 4,虚拟主机 412 可以接收指示来自 VHD 存储服务器 410 的虚拟硬盘文件集合的副本的一个或多个信息分组。此示例中的集合可包括但不限于用户差异虚拟硬盘文件、机器差异虚拟硬盘文件,和 / 或主虚拟硬盘文件。在此示例中,当客操作系统正在运行时,VHD 存储服务器 410 可以同时地将集合拷贝到本地存储 420。即,在下载集合的副本时,虚拟机 416 可以执行存储在 VHD 存储服务 410 中的集合的客操作系统。

[0072] 转回到图 10,操作 1006 示出在实施例,虚拟主机可包括响应于确定虚拟硬盘文件集合的副本已被存储到附连于第一服务器的存储设备中,执行虚拟硬盘文件集合的副本的客操作系统的线路。转回到图 4,在示例性实施例中,虚拟硬盘文件集合可由虚拟主机 412 成功地接收并存储在本地存储 420 中。在此示例中,虚拟平台 402 可接收指示存储操作已结束的信号,并且设定配置文件和 / 或虚拟硬盘文件中的父位置信息,以便位置信息指向集合的本地副本。在修改父位置信息之后,虚拟平台 402 可以使用已更新的位置信息向集合的本地副本发出 IO 任务。因此,由于虚拟平台 402 不再需要通过内部网络从 VHD 存储服务 410 读取数据,因此减轻了 VHD 存储服务 410 的处理压力以及对带宽的需求。

[0073] 参照图 11,其示出图 10 的操作过程的替换实施例,包括操作 1108-1122。转向操作 1108,它示出在实施例,虚拟主机 412 可包括用于从虚拟硬盘文件集合中引导客操作系统的线路,其中虚拟硬盘文件集合包括主虚拟硬盘文件、机器差异虚拟硬盘文件和用户差异虚拟硬盘文件。例如,转向图 5,在此示例性实施例中,虚拟主机 412 可以从 VHD 存储服务器 410 引导客操作系统。在此示例中,存储其上的 VHD 文件集合可包括主 VHD 文件 504、机器差异 VHD 文件 506 和用户差异 VHD 文件,主 VHD 文件可包括操作系统的副本,机器差异 VHD 文件可包括为诸如机器标识符等的特定用户定制操作系统的信息,用户差异 VHD 文件可包括在会话期间已被写到磁盘的数据。

[0074] 在包括操作 1108 的示例性实施例中,用于 VHD 文件集合的拷贝操作可在引导客操作系统之前已经开始,并且可以在客操作系统开始执行之后结束,即,在引导加载器退出并将对虚拟机 416 的控制转移给客 OS 之后。当虚拟平台 402 接收请求以引导客 OS 时,VHD A 的全部文件已被存储到 VHD 存储服务器 410 中,并且虚拟平台 402 可从配置文件确定用户差异 VHD 文件 508 的位置。虚拟平台 402 可以使用用户差异 VHD 文件 508 中的信息来确定机器差异 VHD 文件 506 的位置,并使用机器差异 VHD 文件 506 中的位置信息来确定主 VHD 文件 504 的位置。在下载集合之后,虚拟平台 402 可以将配置文件中的位置设为指向本地用户差异 VHD 文件 514 的位置,将本地用户差异 VHD 文件 514 中的父位置信息设为指向本地机器差异 VHD 文件 512,并将本地机器差异 VHD 文件 512 中的父信息设为指向本地主 VHD 文件 510。在此处,虚拟平台 402 将使用 VHD 文件的本地副本用于 IO。

[0075] 继续对图 11 的描述,操作 1110 示出在实施例,虚拟主机 412 可包括用于从虚拟硬盘文件集合以及存储在附连于第一服务器的存储设备中的主虚拟硬盘文件的预先高速缓存的副本引导客操作系统的线路。例如,转向图 6,在示例性实施例中,集合可包括机器差异 VHD 文件 506 和用户差异 VHD 文件 508。主 VHD 文件 504 的副本可以作为本地主 VHD 文件 510 被预先高速缓存到本地存储 420 中。在此示例中,可以在数据中心中使用存储在本地存储 420 的本地主 VHD 文件 510 来部署虚拟主机 412,或者当开启虚拟主机 412 时,虚拟

主机可以预先拷贝主 VHD 文件 504。在此配置中,消除了由迁移主 VHD 文件 504 引起的可达 8GB 的网络通信。此外,可以本地处理定向到主 VHD 文件中的块的任意读 IO 任务。

[0076] 继续对图 11 的描述,操作 1112 示出在实施例中,虚拟主机 412 可包括用于从虚拟硬盘文件集合以及存储在附连于第一服务器的存储设备中的用户差异虚拟硬盘文件引导客操作系统的线路。在此示例性实施例中,转向图 7,集合可包括主 VHD 文件 504 和机器差异 VHD 文件 506。当客操作系统正在引导时,即引导加载器控制虚拟机 416 时,虚拟平台 402 可通过从本地用户差异 VHD 文件 514 读和写来处理某些 IO 任务,并且通过从存储在 VHD 存储服务器 410 的集合中读取来处理其他任务。在此示例中,在引导之前本地用户差异 VHD 文件 514 可以为空,或者,当虚拟平台 402 接收请求以实例化虚拟机 416 时可以动态地创建本地用户差异 VHD 文件。在第一种情形中,配置文件可包括空的本地用户差异 VHD 文件 514 的位置。在第二种中,配置文件可包括机器差异 VHD 文件 506 的位置,并且虚拟平台 402 可以动态地生成本地用户差异 VHD 文件 514 并将其链接到机器差异 VHD 文件 506。

[0077] 在任一情形中,可以确定机器差异 VHD 文件 506 和主 VHD 文件 504 的位置,并且存储传送引擎 424 可以开始将其拷贝到本地存储 420,从而存储本地机器差异 VHD 文件 512 和本地主 VHD 文件 510。在下载集合之后,虚拟平台 402 可以修改本地用户差异 VHD 文件 514 中的父位置字段以指向本地机器差异 VHD 文件 512 的位置,并且修改本地机器差异 VHD 文件 512 中的父位置字段以指向本地主 VHD 文件 510 的位置。在此处,虚拟平台 402 将使用 VHD 文件的本地副本用于 IO。

[0078] 继续对图 11 的描述,操作 1114 示出在实施例中,虚拟主机 412 可包括被配置为将用户差异虚拟硬盘文件合并到机器差异虚拟硬盘文件中并将机器差异虚拟硬盘文件发送到第二服务器的线路。在退出后持久存储用户数据的部署中,虚拟平台 402 可被配置为将本地用户差异文件 514 中的数据与本地机器差异 VHD 文件 512 中的数据相合并,并将已修改的机器差异 VHD 文件发送回到 VHD 存储服务器 410。例如,虚拟平台 402 可以通过将本地用户差异文件 514 中的每个数据块拷贝到本地机器差异 VHD 文件 512 的对应块中来合并两个文件。在 VHD 存储服务器 410 接收本地机器差异 VHD 文件 512 的已修改的副本之后,VHD 存储服务器将其存储为机器差异 VHD 文件 506。

[0079] 在已修改的版本被设为机器差异 VHD 文件 506 之后,虚拟平台 402 可以可选地删除本地用户差异 VHD 文件 514 和本地机器差异 VHD 文件 512。例如,虚拟平台 402 可被配置为在限制本地存储 420 的存储容量的实例中删除这些文件。在虚拟主机 412 将机器差异 VHD 文件的已合并副本作为本地机器差异 VHD 文件 512 持久存储在本地存储 420 的实例中,虚拟主机 412 可以发送信号到管理员服务器 408 以通知它,当客户机 406 下次连接时,虚拟主机 412 应当托管虚拟机 416。

[0080] 继续对图 11 的描述,操作 1116 示出在实施例中,存储服务器 410 可包括用于将从第一服务器接收的用户差异虚拟硬盘文件合并到存储在第二服务器的机器差异虚拟硬盘文件中的线路。在此示例性实施例中,当用户退出后,虚拟平台 402 可以将本地用户差异文件 514 发送回到 VHD 存储服务器 410 并且指示存储传送引擎 424 合并本地用户差异文件 514 与机器差异 VHD 文件 506。VHD 存储服务器 410 随后可以将已修改的机器差异 VHD 文件存储为机器差异 VHD 文件 506。

[0081] 继续对图 11 的描述,操作 1118 示出在实施例中,存储服务器可被配置为从第一服

务器接收至少用户差异虚拟硬盘文件；并且将至少用户差异虚拟硬盘文件发送到第三服务器。例如，转向图 9，在示例性实施例中，实况迁移操作可用于当虚拟机 416 正在运行时将虚拟机 416 从虚拟主机 412 迁移到虚拟主机 414。例如，存储传送引擎 424 可以接收指示虚拟机 416 将要从虚拟主机 412 迁移到虚拟主机 414 的信号。作为响应，存储传送引擎 424 可以确定需要拷贝哪些虚拟硬盘文件以在虚拟主机 414 上运行虚拟机 416。在此示例中，存储传送引擎 424 可以确定需要拷贝至少本地用户差异 VHD 文件 514。在此示例中，可以使用此处描述的镜像技术将本地用户差异 VHD 文件 514 发送到 VHD 存储服务 410。

[0082] 发送存储在用户差异 VHD 文件 508 中的至少信息的操作与将虚拟机 416 的状态（例如，RAM 的内容和虚拟设备的状态）从虚拟主机 412 拷贝到虚拟主机 414 的操作可以并行进行。在另一个配置中，可在实况迁移操作开始之前进行存储拷贝操作。在任一情形中，可以将只读文件拷贝到 VHD 存储服务器 410 并且可以镜像读 / 写文件。例如，运行在虚拟主机 412 上的存储传送引擎 424 可以顺序地将来自本地用户差异 VHD 文件 514 的数据块拷贝到 VHD 存储服务器 410，而将写发送到本地用户差异 VHD 文件 514 和用户差异 VHD 文件 508 二者。当拷贝操作结束时，虚拟平台 402 可以修改配置文件以将位置信息设为指向用户差异 VHD 文件 508 的位置，从而将写发送到 VHD 存储服务器 410。当对虚拟机 416 的控制从虚拟平台 402 切换到虚拟平台 404 之后，虚拟平台 404 可以通过内部网络从 VHD 存储服务器 410 读取数据块。在此处，存储传送引擎 424 可被配置为将至少用户差异 VHD 文件 508 拷贝到虚拟主机 414。

[0083] 继续对图 11 的描述，操作 1120 示出在实施例中，存储服务器 410 可包括用于将虚拟硬盘文件集合的至少部分预先高速缓存到第二服务器的随机访问存储器中的线路。例如，转回到图 5，在示例性实施例中，作出关于在一种类型的操作系统的引导操作期间通常读取什么字节的判断。例如，可以通过连续的引导操作跟踪由虚拟平台 402 发出的读 I/O 任务，并且确定通常读取的一组字节。在此示例中，VHD 存储服务器 410 可被配置为在其引导时将这些字节预先高速缓存到 RAM 中，以便具有用于引导虚拟机的可用字节。这可以加快虚拟机的引导操作。例如，用于存储虚拟硬盘驱动器的持久存储可以是商用硬件，例如现货供应的 SATA 驱动器，并且使 VHD 存储服务器 410 读取这种类型的硬件的字节带来可由通过将字节预先高速缓存到 RAM 消除的另一层等待时间。

[0084] 继续对图 11 的描述，操作 1122 示出在实施例中，管理员服务器可包括用于响应于确定第一服务器包括机器差异虚拟硬盘文件的当前副本，指示客户机计算机系统连接到第一服务器的线路。例如，转回到图 4，在实施例中，管理员服务器 408 可以包括线路，例如，由指示连接中介 426 的可执行指令配置的处理器。在此示例中，指令可以配置处理器以管理信息表，该信息表标识哪个虚拟主机具有带有指示其上一次修改时间的时戳的 VHD 文件。响应于从客户机 406 接收连接到虚拟机 416 的连接请求，管理员服务器 408 可以执行程序，该程序检查 VHD 存储服务器 410 中虚拟硬盘 A 的 VHD 文件集合的时戳，并将该时戳与表中的时戳进行比较。在存在匹配的实例中，例如，虚拟主机 412 包括具有与存储在 VHD 存储服务器 410 中的与机器差异虚拟硬盘文件相关联的时戳相同的时戳的机器差异虚拟硬盘文件，管理员服务器 408 可以试图将虚拟机 416 加载到虚拟主机 412 上。这将消除必须将 VHD 文件中的至少某些拷贝到虚拟主机 412。在大多数情形中，如果虚拟主机 412 包括机器差异 VHD 文件的当前副本，则它还包括主 VHD 文件的当前副本。如果情况如此，则不必从 VHD 存

储服务器 410 下载任何文件。

[0085] 转向操作 1124, 它示出响应于确定第一服务器上的负载为低, 指示客户机计算机系统连接到第一服务器。例如, 在实施例中, 管理员服务器 408 可以包括线路, 例如, 由指示连接中介 426 的可执行指令配置的处理器。在此示例中, 该指令可以配置处理器以管理信息表, 该信息表标识每个虚拟服务器当前所处的当前负载。例如, 负载信息可指示分配给虚拟机的 RAM 的数量以及每个虚拟服务器具有的 RAM 的总量。响应于从客户机 406 接收连接到虚拟桌面会话的连接请求, 管理员服务器 408 可以执行程序, 该程序检查每个服务器上的负载并选择最少负载的服务器来托管虚拟桌面会话。

[0086] 在执行操作 1122 和 1124 二者的示例实施例中, 管理员服务器 408 可被配置为标识其上具有低负载的服务器, 并随后检查 VHD 存储服务器 410 中虚拟硬盘 A 的 VHD 文件集合的时戳, 并且将该时戳与表中的时戳进行比较。在存在匹配的实例中, 例如, 虚拟主机 412 包括具有与存储在 VHD 存储服务器 410 中的与机器差异虚拟硬盘文件相关联的时戳相同的时戳的机器差异虚拟硬盘文件, 管理员服务器 408 可以试图将虚拟机 416 加载到虚拟主机 412 上。

[0087] 现在转向图 12, 它示出用于描述涉及从联网存储执行虚拟机的操作的示例性实施例。操作 1200 开始操作过程, 并且操作 1202 示出操作, 从存储在第二服务器中的机器差异虚拟硬盘文件和存储在附连于第一服务器的存储设备中的主虚拟硬盘文件引导操作系统, 其中主虚拟硬盘文件包括操作系统的副本, 并且机器差异虚拟硬盘文件包括用于为用户定制操作系统的副本的机器标识信息。例如, 参考图 6 或 8, 在示例性实施例中, 虚拟主机 412 可包括本地主 VHD 文件 510。例如, 本地主 VHD 文件 510 可以之前已经被下载, 或者可以使用副本来部署虚拟主机 412。在任一情形中, 虚拟平台 402 可从本地主 VHD 文件 510 中的信息的组合和机器差异 VHD 文件 506 来引导客操作系统。

[0088] 虚拟平台 402 可以响应于从图 4 的管理员服务器 408 接收的请求来启动引导操作。虚拟平台 402 可以接收与经由图 4 的客户机 406 登录的用户的用户账户相关联的配置文件。虚拟平台 402 可以访问配置文件, 并且存储其中的字段可包括例如用户差异 VHD 磁盘文件的位置, 用户差异 VHD 磁盘文件可被存储到本地存储 420 或 VHD 存储服务器 410。根据用户差异 VHD 磁盘文件及其父的位置, 可以发现机器差异 VHD 文件 506。在此示例中, 机器差异 VHD 文件 506 可位于 VHD 存储服务器 410 的存储中, 并且其父可以是本地主 VHD 文件 510。在打开构成 VHD A 的 VHD 文件之后, 虚拟平台 402 可以通过将由机器差异 VHD 文件 506 处理的读 IO 请求发送到 VHD 存储服务器 410 并将由本地主 VHD 文件 510 处理的读 IO 请求发送到本地存储 420 来开始引导客操作系统。

[0089] 继续对图 12 的描述, 操作 1204 示出了用于在同时引导客操作系统时, 从第二服务器下载机器差异虚拟硬盘文件的副本的操作。例如, 转向图 6 或 8, 虚拟主机 412 可从 VHD 存储服务器 410 接收指示机器差异 VHD 文件 506 的副本的一个或多个信息分组, 并且将副本存储为本地机器差异 VHD 文件 512。在此示例中, VHD 存储服务器 410 可以在引导操作期间同时地将机器差异 VHD 文件 506 拷贝到本地存储 420。即, 当虚拟机 416 从组合进行引导时, 拷贝正同时地被下载。

[0090] 再次转回到图 12, 操作 1206 示出用于响应于确定机器差异虚拟硬盘文件已被下载, 从机器差异虚拟硬盘文件的副本和主虚拟硬盘文件执行客操作系统的操作。转回到图

6 或 8, 在示例性实施例中, 本地机器差异 VHD 文件 512 可被成功地存储到虚拟主机 412 的文件系统中。在此示例中, 虚拟平台 402 可接收指示存储操作已结束的信号, 并且将用户差异 VHD 文件的父位置字段中的父位置信息设为指向本地机器差异 VHD 文件 512 的位置。在修改位置之后, 虚拟平台 402 可以使用已更新的位置信息来执行对于本地副本的 IO 操作。

[0091] 现在转向图 13, 其示出图 12 所描绘的操作过程的替换实施例, 包括操作 1308-1320。如操作 1308 所示, 在实施例中, 虚拟主机 412 可包括用于从第二服务器下载用户差异虚拟硬盘文件的线路, 用户差异虚拟硬盘文件存储在会话期间写到虚拟硬盘的数据块。转向图 6, 在此实施例中, 用户差异 VHD 磁盘 508 可从 VHD 存储服务器 410 拷贝到虚拟主机 412, 并被存储为本地用户差异 VHD 磁盘 514。如图所示, 在此示例中, 虚拟平台 402 可以接收包括父位置字段的配置文件, 该父位置字段包括用户差异 VHD 文件 508 的位置。用户差异 VHD 文件 508 中的父位置字段可包括机器差异 VHD 文件 506 的位置, 并且其父字段可包括本地主 VHD 文件 510 的位置。存储传送引擎 424 可以顺序地将来自用户差异 VHD 文件 508 的数据块拷贝到本地用户差异 VHD 文件 514。当数据块被拷贝时, 虚拟平台 402 可向本地用户差异 VHD 文件 514 和用户差异 VHD 文件 518 二者发送写, 直到用户差异 VHD 文件 508 的最后一块被拷贝。在此处, 虚拟平台 402 可被配置为停止向用户差异 VHD 文件 508 写入, 并可修改配置文件中的位置字段以指向本地用户差异 VHD 文件 514 的位置而非用户差异 VHD 文件 508。在此处, 虚拟平台 402 将使用用户差异 VHD 文件 514 的本地副本用于相关联的 IO 操作。

[0092] 继续对图 13 的描述, 它示出了操作 1310, 该操作是用于将数据块写入被预先高速缓存到附连于第一服务器的存储设备中的用户差异虚拟硬盘文件的操作。例如, 转向图 8, 在此示例性实施例中, 本地存储 420 可包括本地用户差异 VHD 文件 514。在此示例性实施例中, 用户差异 VHD 文件 514 可被预先高速缓存到本地存储 420, 从而不必跨网络对其进行拷贝。例如, 当虚拟平台 402 引导虚拟机 416 时, 其可被配置为生成空的本地用户差异 VHD 文件 514。虚拟处理器向 VHD A 写入时, 该写可被存储到本地用户差异 VHD 文件 514 的块中, 并且可以从 0 字节增长。在此示例性实施例中, 每次虚拟平台 402 向虚拟平台 402 写入时, 它可以在位图中设置一位。当由虚拟平台 402 接收读 IO 任务时, 其可被配置为检查位图以查看是否为该块设置了位。若是, 则虚拟平台 402 从本地用户差异 VHD 文件 514 读取存储在块中的字节。在未设置位的情况下, 虚拟平台 402 可被配置为检查用于本地机器差异 VHD 文件 512 的位图。如果用于本地机器差异 VHD 文件 512 的位图未被设置, 则虚拟平台 402 可从本地主 VHD 文件 510 读取块。

[0093] 操作 1312 示出在实施例中, 虚拟主机 412 可包括用于将用户差异虚拟硬盘文件与机器差异虚拟硬盘文件的副本进行合并, 并将机器差异虚拟硬盘文件的副本发送到第二服务器的线路。在持久存储用户数据的部署中, 虚拟平台 402 可被配置为将本地用户差异文件 514 中的数据与本地机器差异 VHD 文件 512 中的数据进行合并, 并将已修改的机器差异 VHD 文件发送回到 VHD 存储服务器 410。例如, 虚拟平台 402 可以通过将本地用户差异文件 514 中的每个数据块拷贝到本地机器差异 VHD 文件 512 的对应块中来合并两个文件。在 VHD 存储服务器 410 接收本地机器差异 VHD 文件 512 的已修改的副本之后, VHD 存储服务器将其存储为机器差异 VHD 文件 506。

[0094] 操作 1314 示出在实施例中, 虚拟主机 412 可包括用于将用户差异虚拟硬盘文件发

送到第二服务器的线路。在此示例性实施例中,当用户退出后,虚拟平台 402 可以将本地用户差异文件 514 发送回到 VHD 存储服务器 410 并且指示存储传送引擎 424 合并本地用户差异文件 514 与机器差异 VHD 文件 506。VHD 存储服务器 410 随后可以将已修改的机器差异 VHD 文件存储为机器差异 VHD 文件 506。

[0095] 操作 1316 示出在实施例中,虚拟主机 412 可包括用于将存储在用户差异虚拟硬盘文件中的数据写入存储在第二服务器中的机器差异虚拟硬盘文件的线路。在此示例性实施例中,在虚拟主机 412 中执行的存储传送引擎 424 可被配置为将本地用户差异 VHD 文件 514 中的差异合并到机器差异 VHD 文件 506 中。例如,存储传送引擎 424 可以打开机器差异 VHD 文件 506,并且将来自用户差异 VHD 文件 514 的数据块跨内部网络写入到机器差异 VHD 文件 506 的对应块中。

[0096] 转向操作 1318,它示出虚拟主机 412 可包括用于在迁移操作期间将用户差异虚拟硬盘文件发送到第二服务器的线路。例如,转向图 9,在示例性实施例中,实况迁移操作可用于当虚拟机 416 正在运行时将虚拟机 416 从虚拟主机 412 迁移到虚拟主机 414。例如,存储传送引擎 424 可以接收指示虚拟机 416 将要从虚拟主机 412 迁移到虚拟主机 414 的信号。作为响应,存储传送引擎 424 可以确定需要拷贝哪些虚拟硬盘文件以便在虚拟主机 414 上运行虚拟机 416。在此示例中,存储传送引擎 424 可以确定需要迁移至少本地用户差异 VHD 文件 514。在此示例中,可以使用此处描述的镜像技术将本地用户差异 VHD 文件 514 发送到 VHD 存储服务 410。此外,该操作可以与将虚拟机 416 的状态从虚拟主机 412 拷贝到虚拟主机 414 的操作并行进行。

[0097] 参考回到图 13,操作 1320 示出使用空闲带宽从第二服务器下载机器差异虚拟硬盘文件的副本。例如,在存储传送引擎 424 执行的实施例中,例如虚拟主机 412 和 / 或 VHD 存储服务器 410 可以确定使用内部网络的空闲带宽来下载机器差异 VHD 文件 506。在此示例中,针对执行虚拟机 416 来优化系统,并且当网络以其他方式空闲时进行拷贝。或换言之,如果带宽的 80% 正被使用,则存储传送引擎 424 可以使用剩余带宽的 20% 将机器差异 VHD 文件 506 发送到虚拟主机 412。

[0098] 现在转向图 14,它示出了操作过程。操作 1400 开始该过程,并且操作 1402 示出用于将客操作系统从虚拟硬盘文件集合引导到虚拟机的操作,其中集合的第一部分被存储到附连于第一服务器的存储设备上,第二部分被存储到第二服务器上。例如,虚拟主机 412 中的处理器可以接收指示虚拟处理器运行引导加载器的线程。处理器可以运行虚拟平台 402 并确定引导加载器正在请求存储在 VHD 文件集合中的字节。VHD 文件集合中的某些可被存储到本地存储 420,并且其余的可被存储到 VHD 存储服务器 410。在此示例中,虚拟平台 420 可以取回字节并将其返回到引导加载器,该引导加载器随后可以继续引导客 OS。

[0099] 继续对图 14 的描述,操作 1404 示出了在引导客操作系统时,下载集合第二部分的至少子部分。例如,当引导加载器正在运行并控制虚拟机 416 时,虚拟主机 412 可以下载指示第一部分的至少子部分的副本的一个或多个信息分组。例如,在存储传送引擎 424 执行的实施例中,例如 VHD 存储服务器 410 可被配置为将其正在存储的虚拟硬盘文件中的至少某些发送到虚拟主机 412。在特定的示例中,转向图 6,存储传送引擎 424 可被配置为将机器差异 VHD 文件 506 和 / 或用户差异 VHD 文件 508 发送到虚拟主机 412。

[0100] 可以基于内部网络的可用带宽和 VHD 存储服务器 410 所处的负载作出关于是否发

送一个或两个的判断。例如,如果 VHD 存储服务器 410 正在处理来自多个虚拟机的多个请求,则其可被配置为延迟迁移 VHD 文件中的某些或全部直到有可用的空闲处理器周期。在相同或另一个实施例中,如果可用带宽为低,则存储传送引擎 424 可被配置为延迟迁移 VHD 文件中的某些或全部直到有可用的带宽。

[0101] 转向操作 1406,它示出响应于确定集合第二部分的至少子部分被下载到附连于第一服务器的存储设备上,执行附连于第一服务器的存储设备上的客操作系统。在示例性实施例中,在引导加载器已经退出之后的某个时刻并且客 OS 受虚拟机 416 的控制,第一部分可被成功地下载到本地存储 420。虚拟平台 402 可以接收指示存储操作已结束的信号,并将任意已下载的 VHD 文件中的父位置字段修改为指向已下载的副本。在此处,虚拟平台 402 可以使用已更新的位置信息来执行针对 VHD 文件的本地副本的 IO 操作。因此,由于虚拟平台 402 不再需要通过内部网络向 VHD 存储服务 410 发送读 / 写操作,因此减轻了 VHD 存储服务 410 的压力以及对带宽的需求。

[0102] 现在参考图 15,它示出图 14 所示的操作过程的替换操作过程,包括细化 1508。如图 15 所示,细化 1508 示出在实施例中,集合的第二部分包括主虚拟硬盘文件和机器差异虚拟硬盘文件,并且集合的第一部分包括用户差异虚拟硬盘文件。在此示例性实施例中,转向图 7,集合可包括主 VHD 文件 504 和机器差异 VHD 文件 506。当客操作系统正在引导时,虚拟平台 402 可以将写 IO 任务和某些读 IO 任务发送到本地用户差异 VHD 文件 514,并将对集合的读 IO 任务发送到 VHD 存储服务器 410。在此示例中,在引导之前本地用户差异 VHD 文件 514 可以为空,或者,当虚拟平台 402 接收请求以实例化虚拟机 416 时可以动态地创建本地用户差异 VHD 文件。在第一种情形中,配置文件可包括空的本地用户差异 VHD 文件 514 的位置。在第二种中,配置文件可包括机器差异 VHD 文件 506 的位置以及指示虚拟平台 402 创建本地用户差异 VHD 文件 514 并将其链接到机器差异 VHD 文件 506 的信息。

[0103] 前述的详细描述通过示例和 / 或操作图阐述了系统和 / 或进程的各实施例。在这样的框图和 / 或示例包含一个或多个功能和 / 或操作的范围内,本领域技术人员将理解,这样的框图,或示例内的每一功能和 / 或操作可以分别地和 / 或共同地通过范围广泛的硬件、软件、固件或几乎其任何组合来实现。

[0104] 尽管示出和描述了此处所描述的本主题的特定方面,但是,对于本领域的技术人员显而易见的是,基于此处的原理,在不偏离此处所描述的主题以及其广泛的方面,可以作出更改和修改,因此,所附权利要求书将包含在它们的范围内,所有这样的更改和修改都将在此处所描述的主题的真正的精神和范围内。

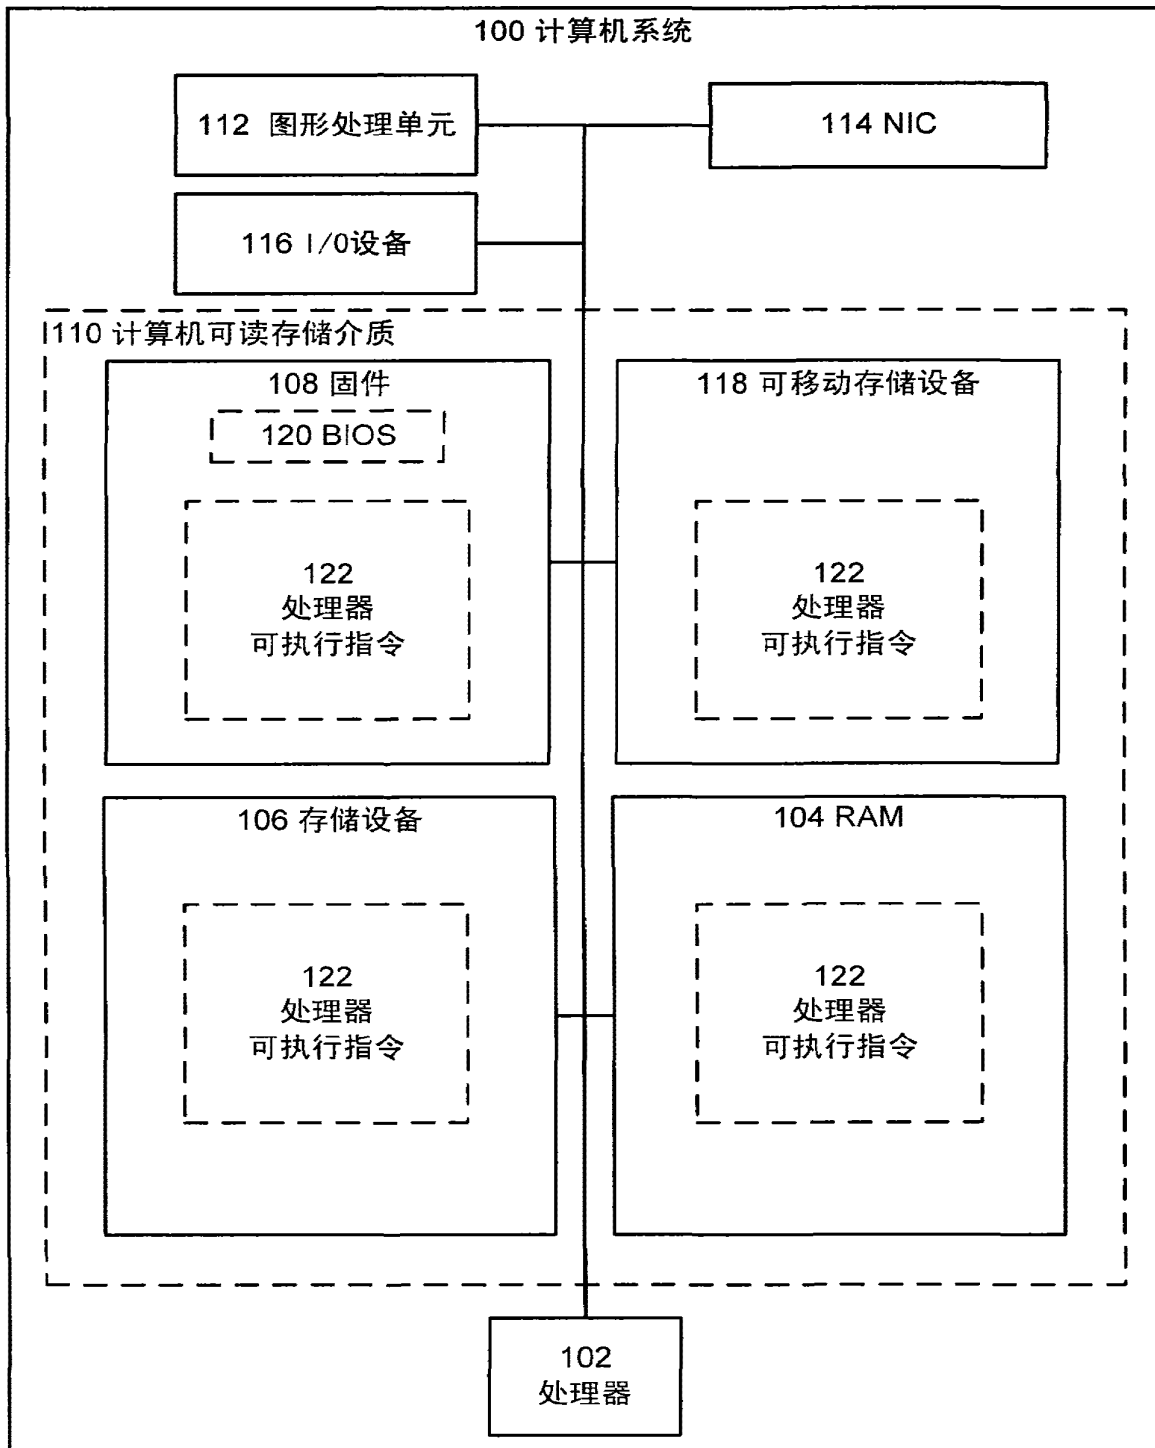


图 1

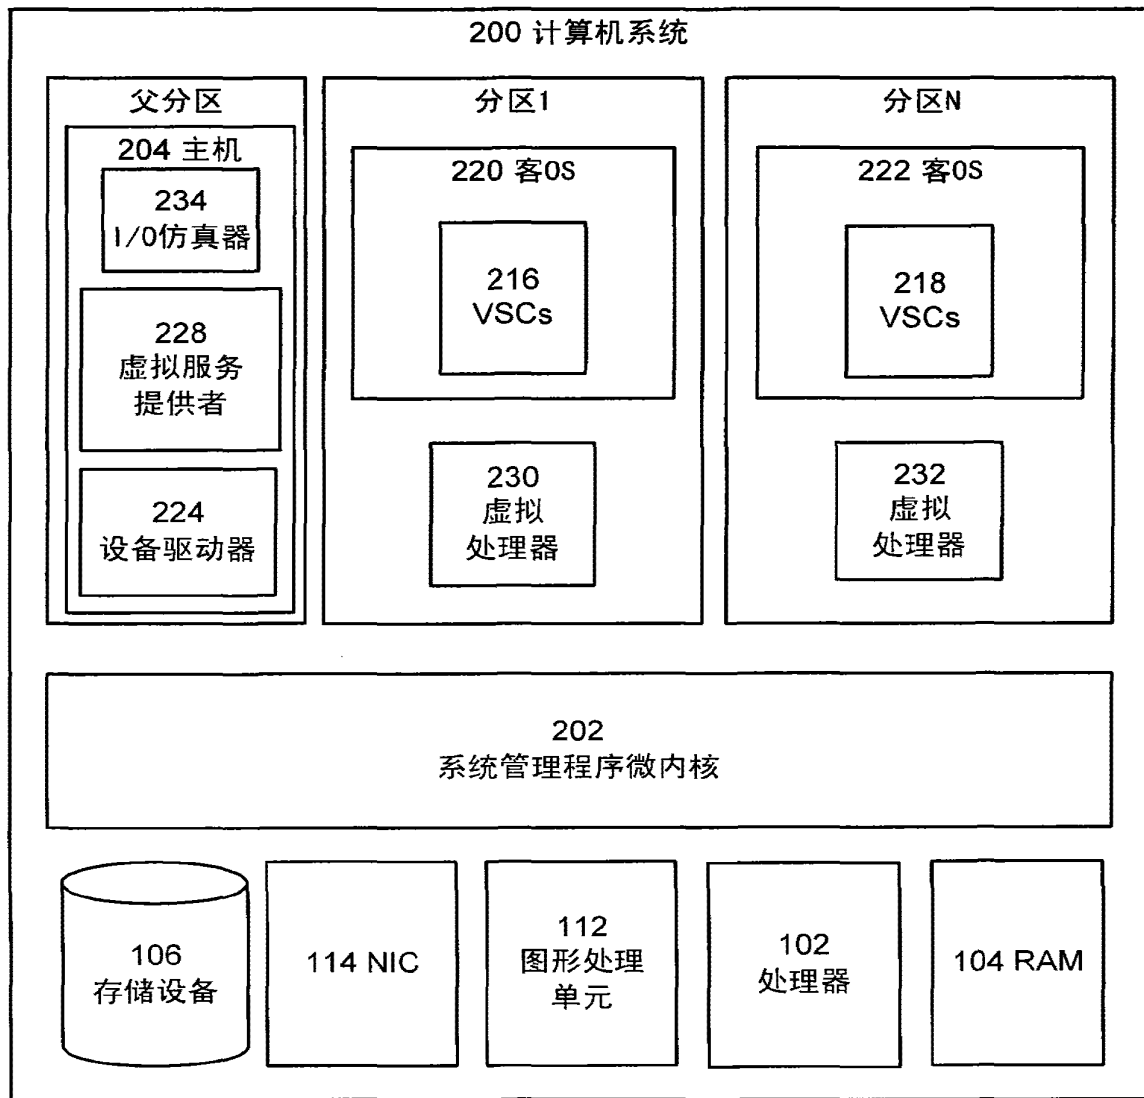


图 2

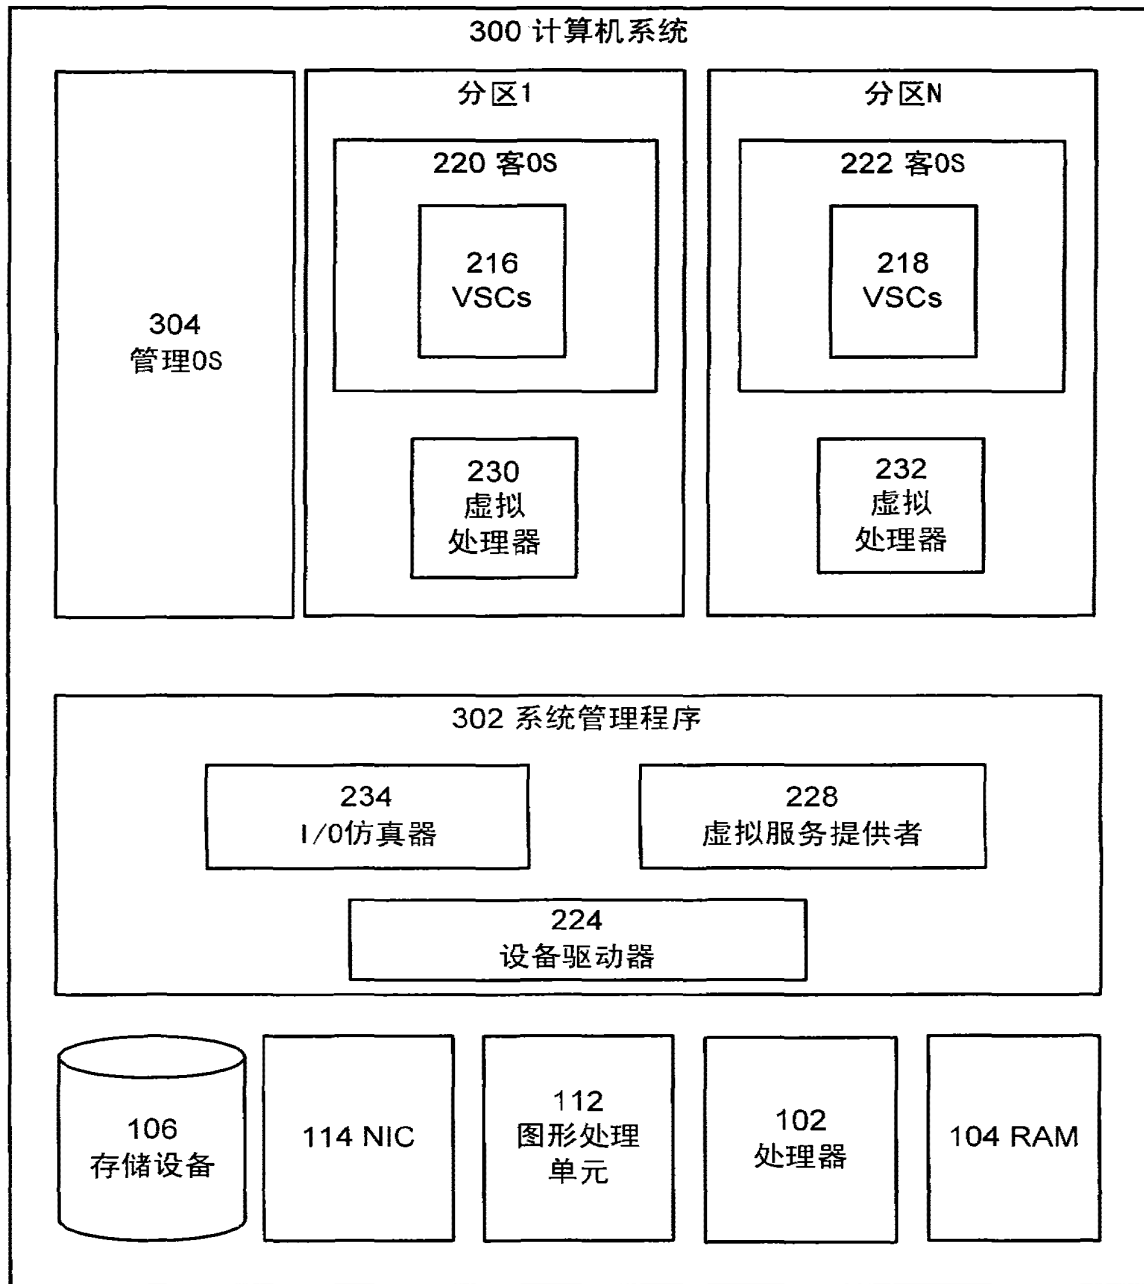


图 3

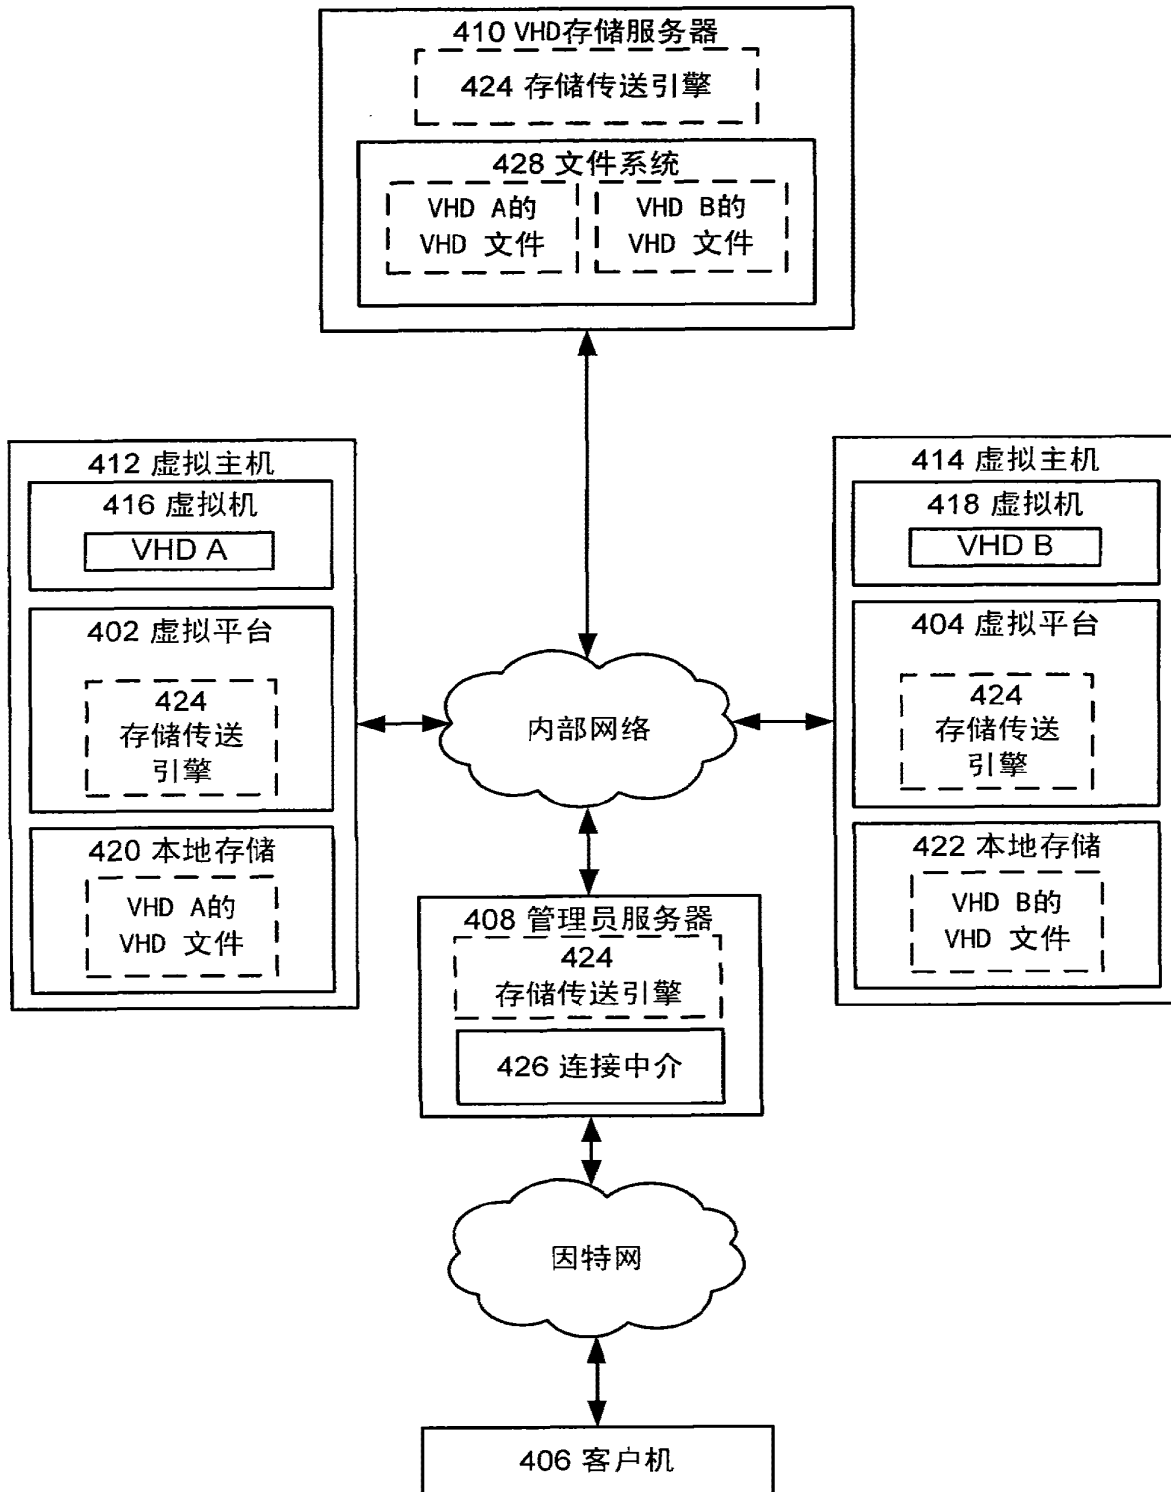


图 4

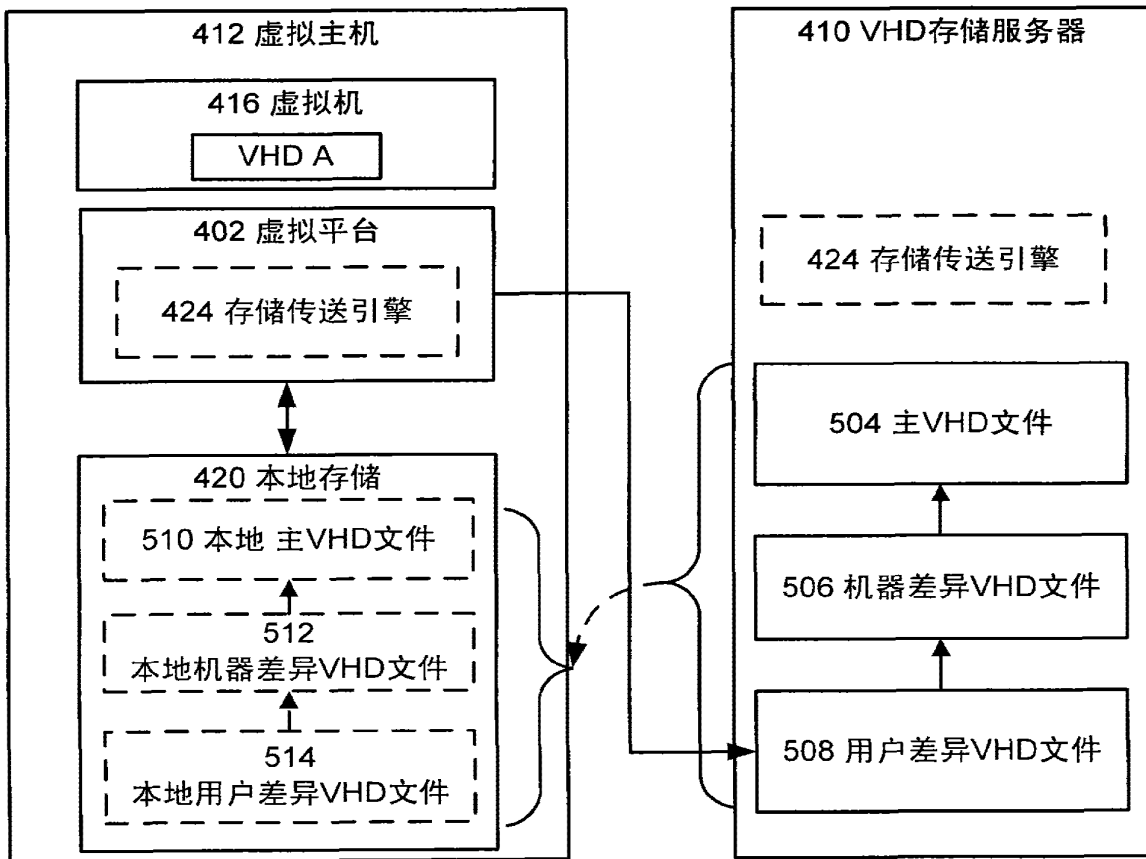


图 5

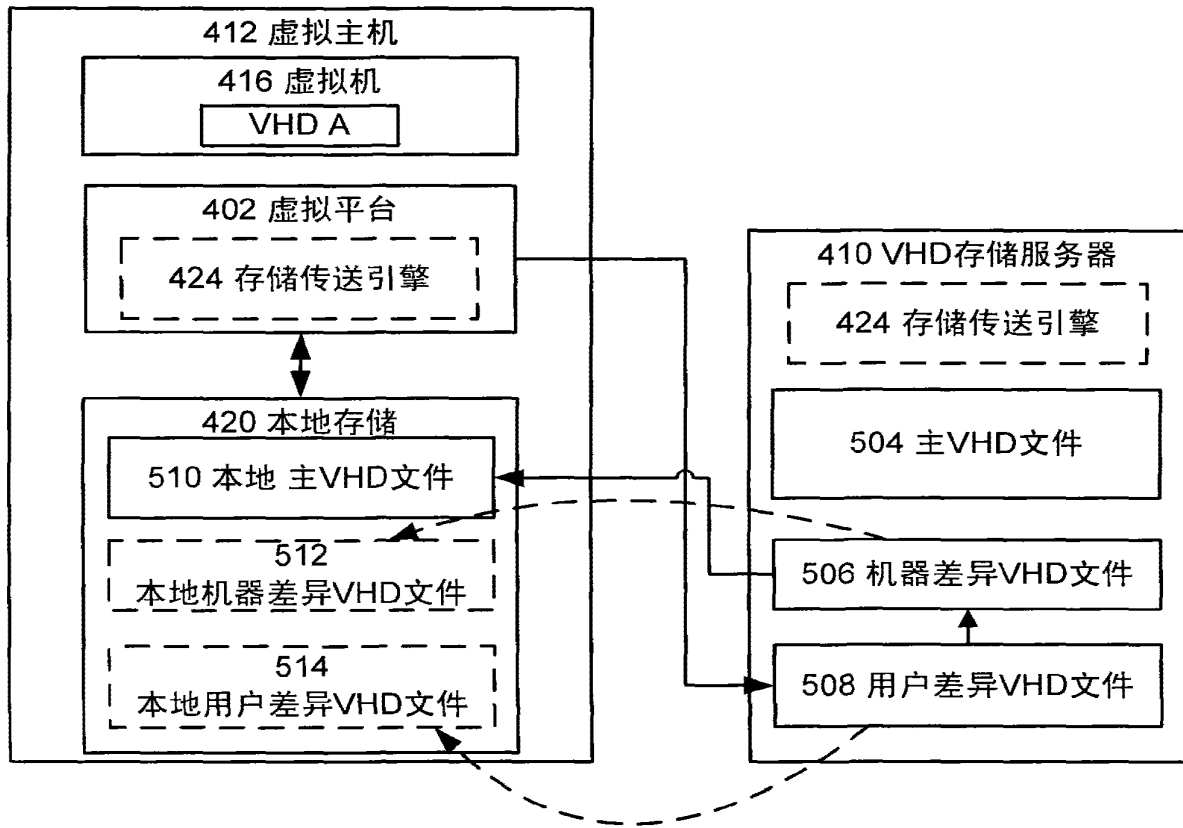


图 6

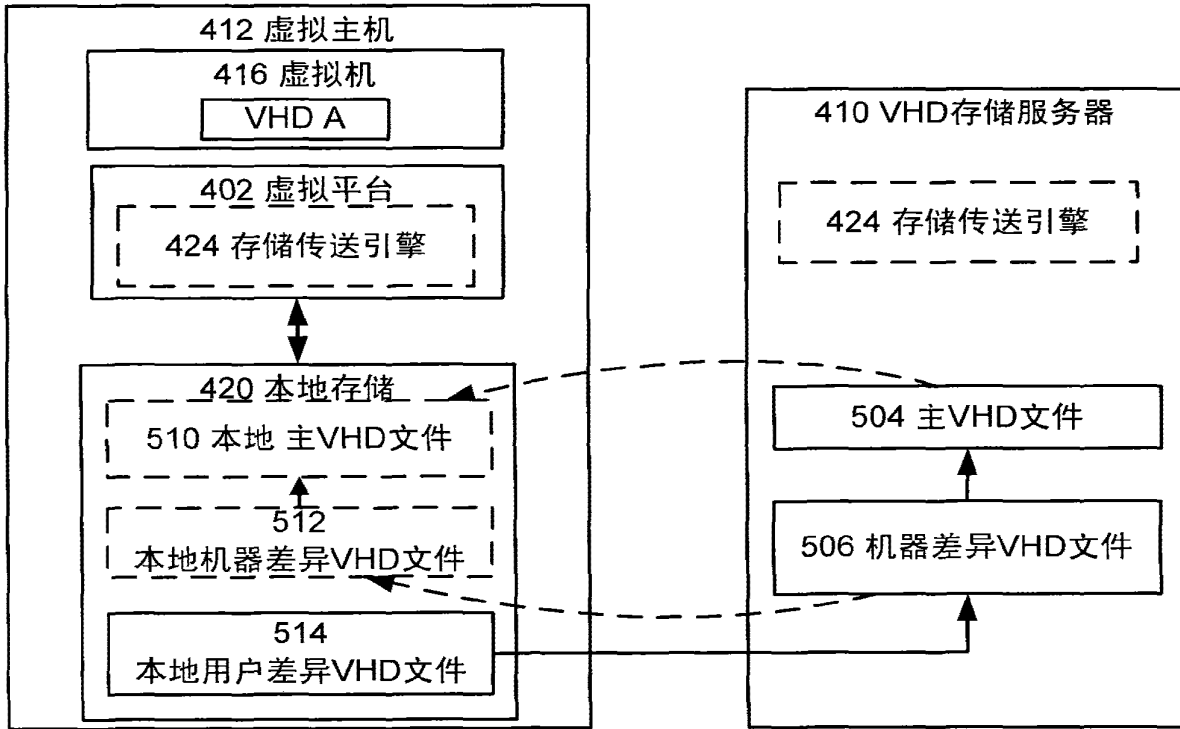


图 7




图 8

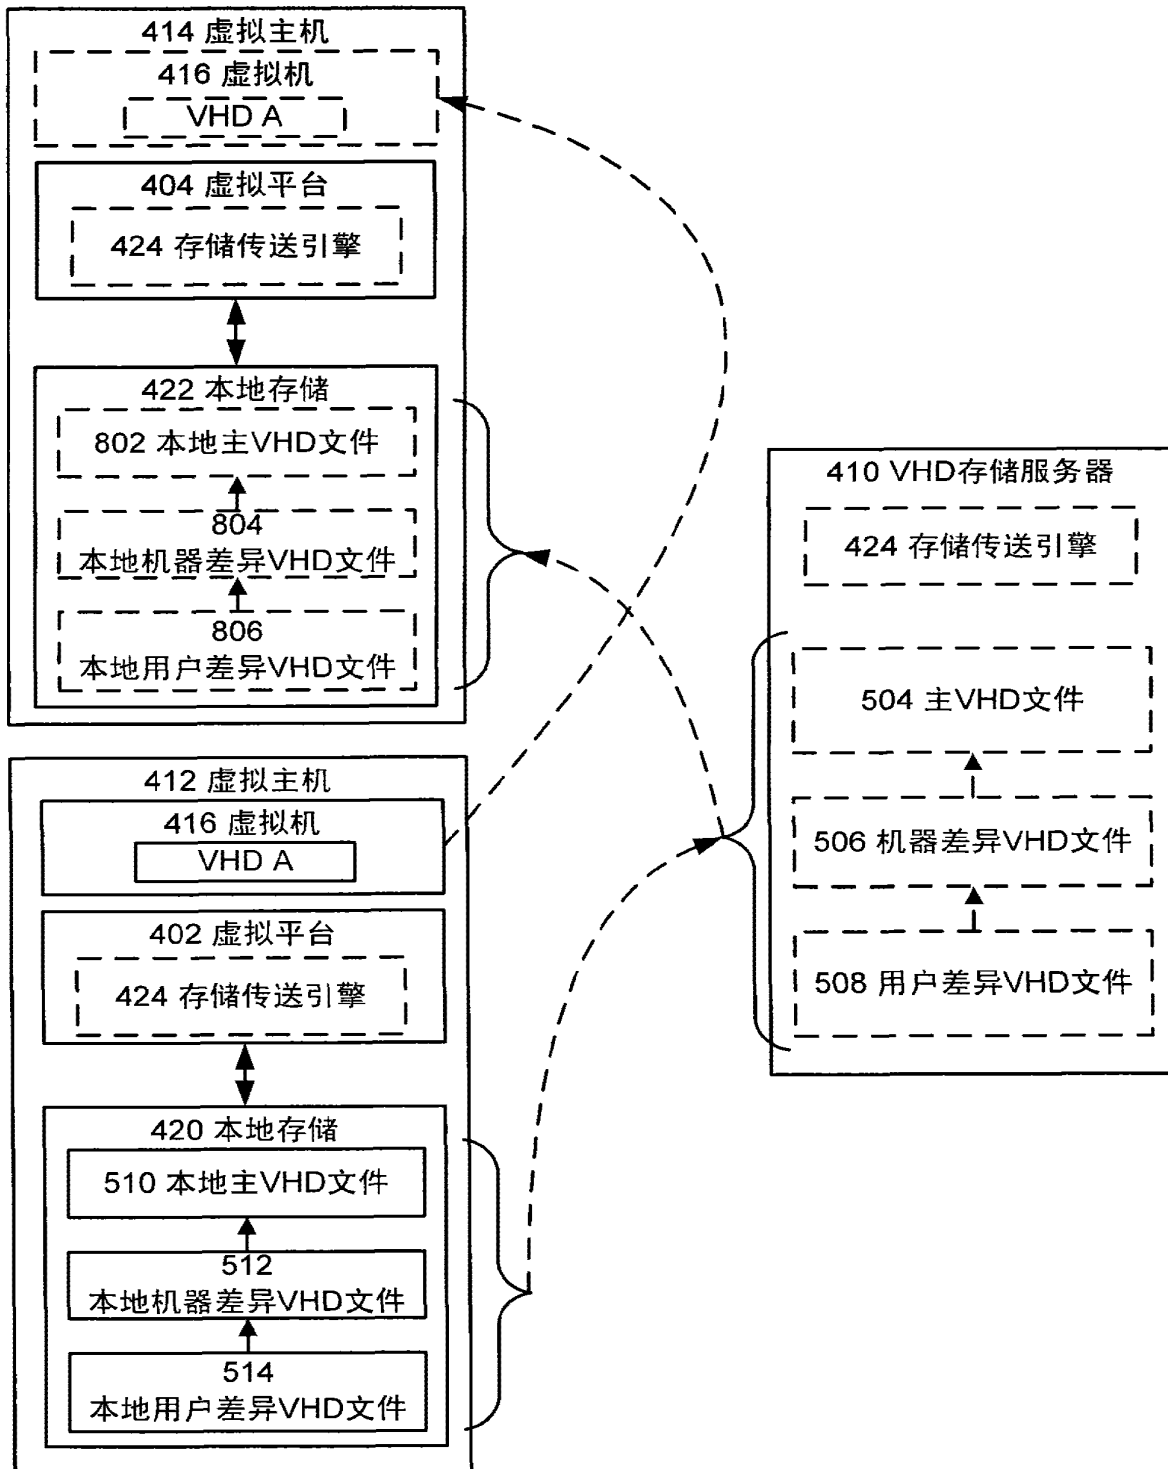


图 9

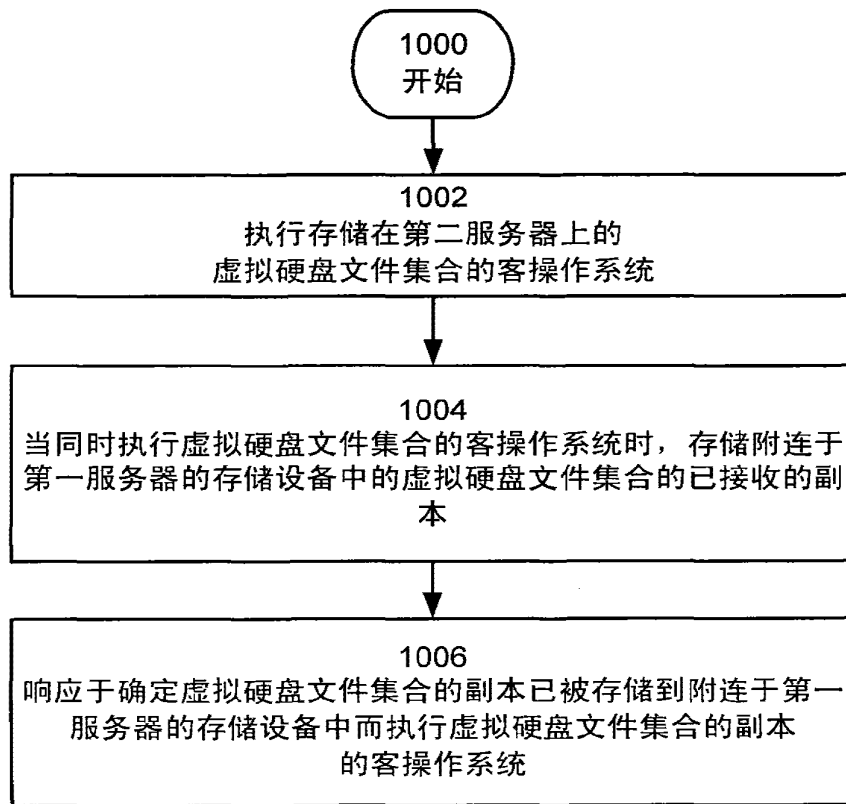


图 10

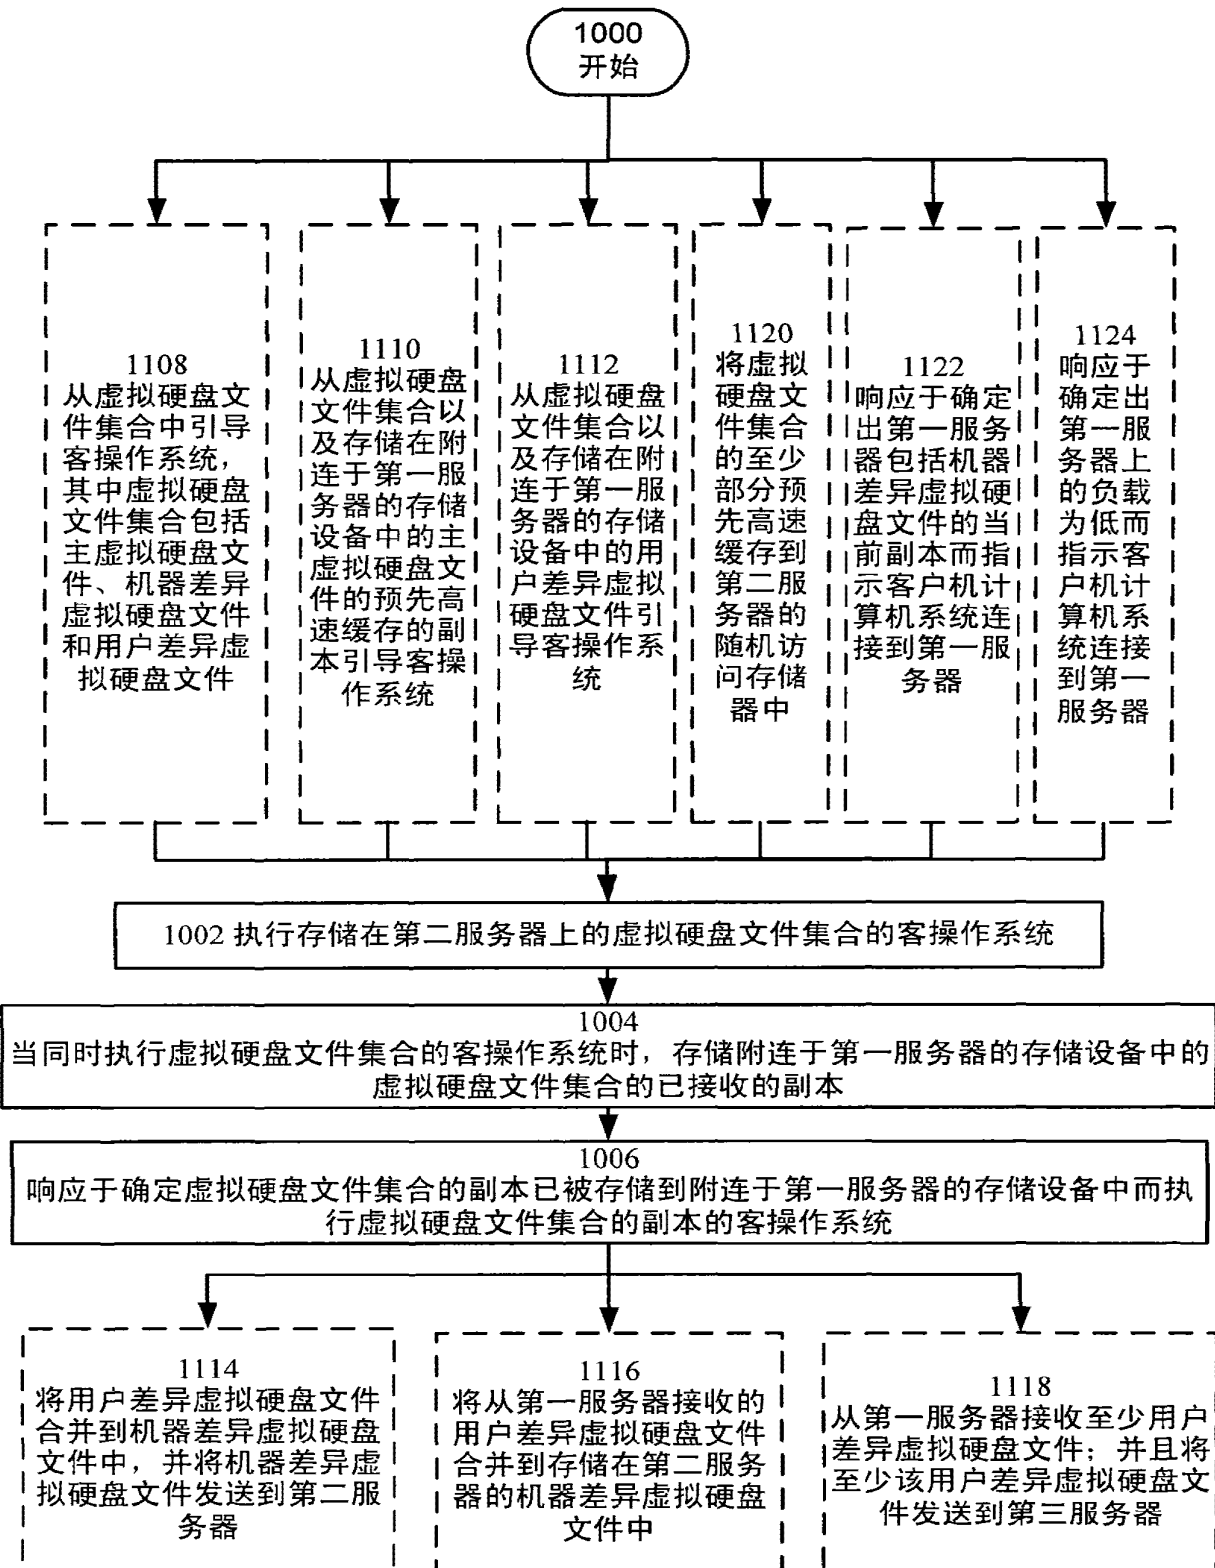


图 11

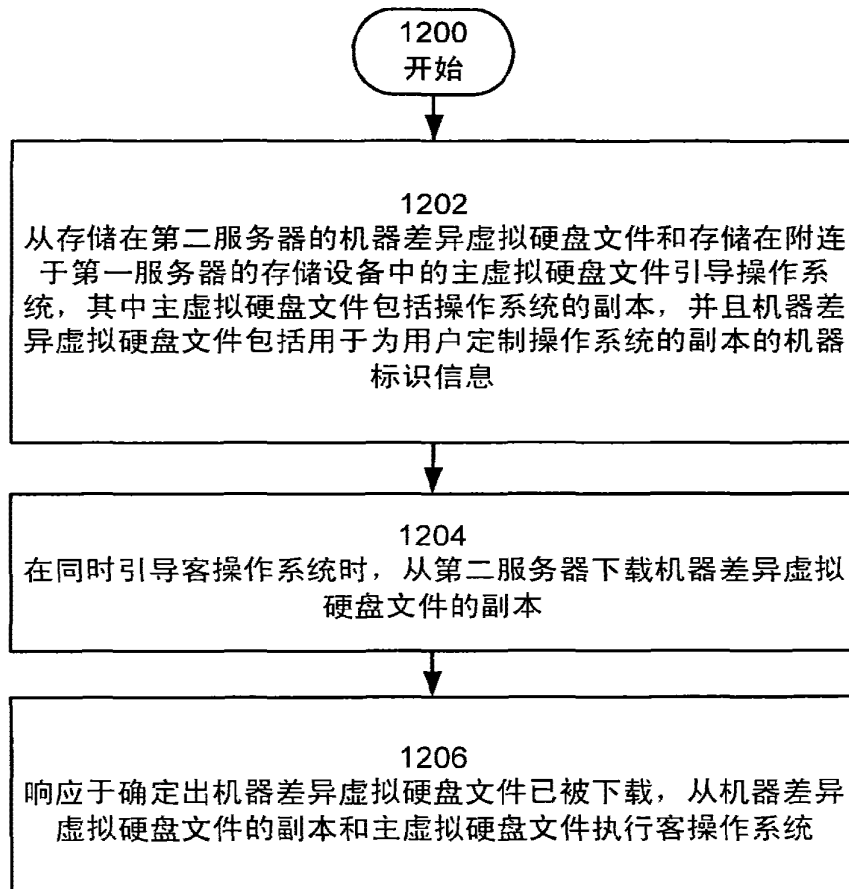


图 12

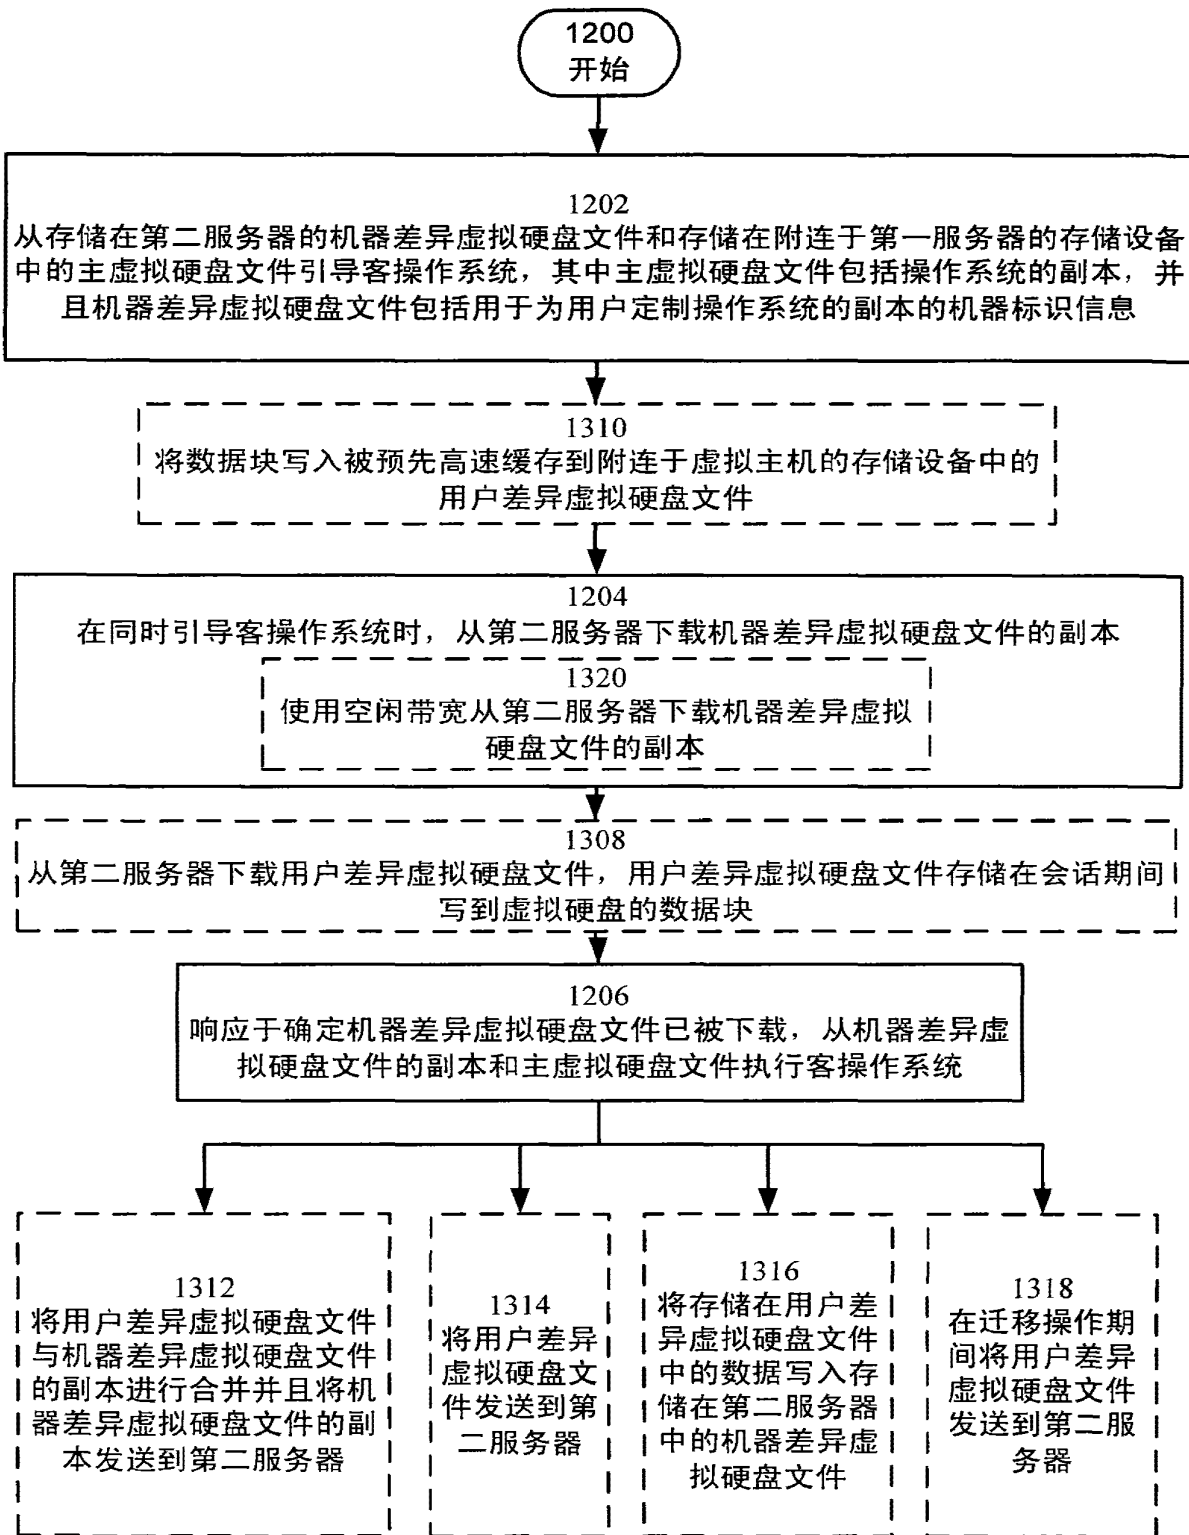


图 13

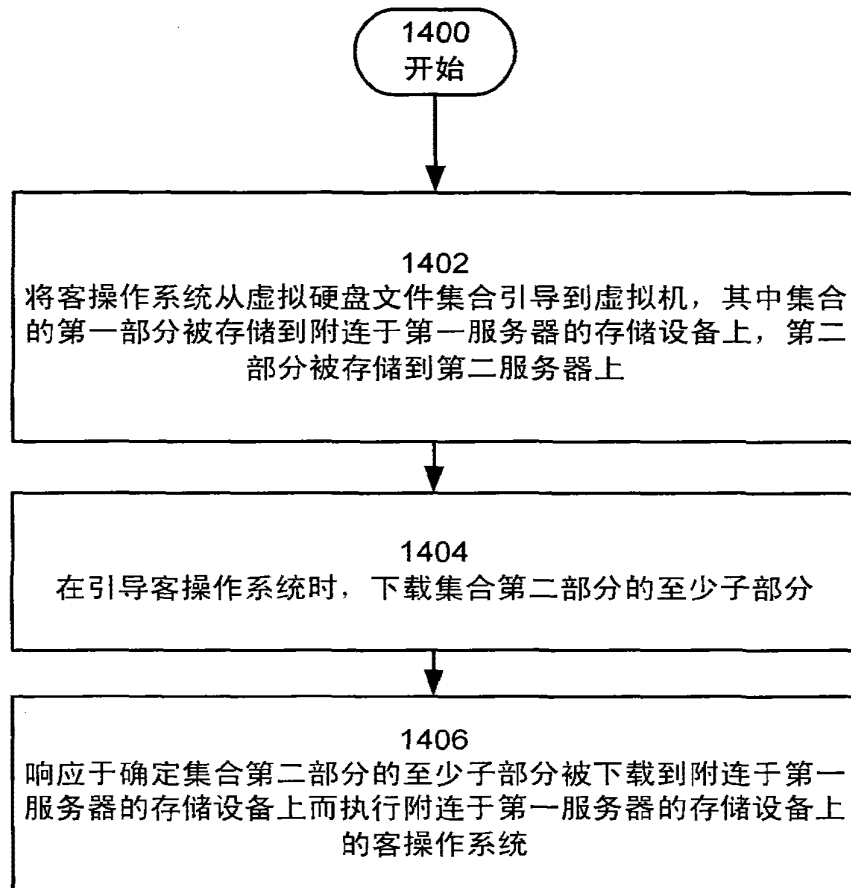


图 14

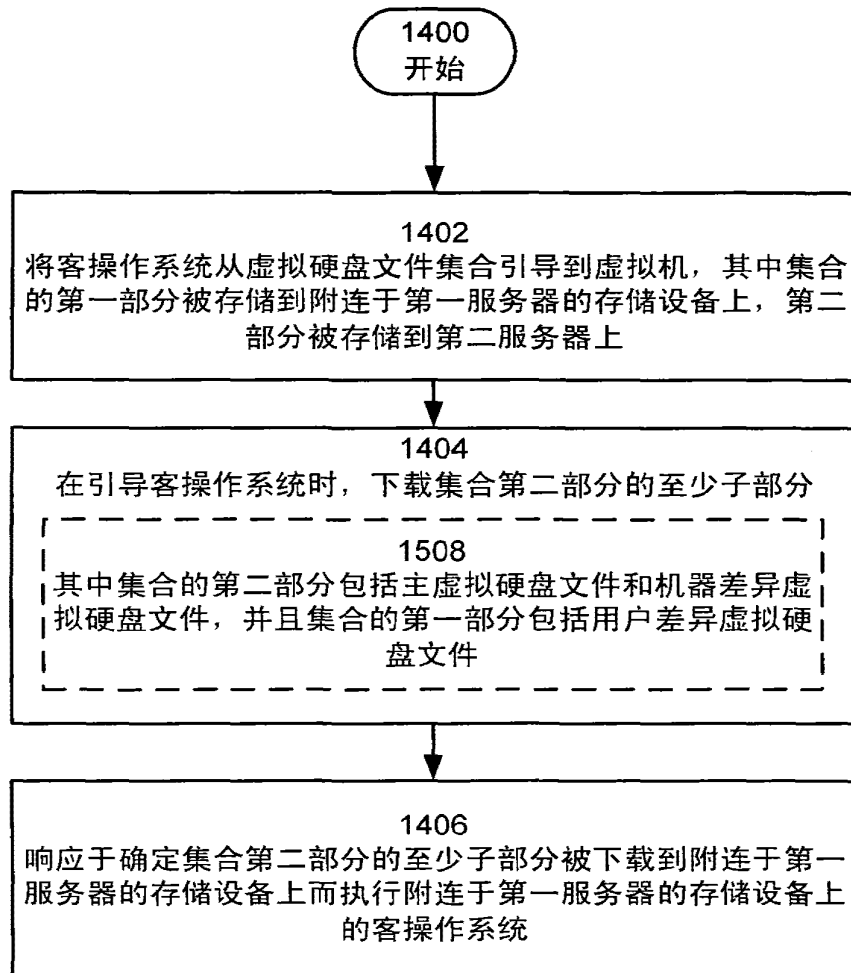


图 15

Electronic Patent Application Fee Transmittal				
Application Number:		15061021		
Filing Date:		04-Mar-2016		
Title of Invention:		MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		
First Named Inventor/Applicant Name:		Matthew A. Markley		
Filer:		Paul J. Otterstedt		
Attorney Docket Number:		YOR920120450US2		
Filed as Large Entity				
Filing Fees for Utility under 35 USC 111(a)				
Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Miscellaneous:				
Submission- Information Disclosure Stmt	1806	1	180	180
Total in USD (\$)				180

Electronic Acknowledgement Receipt

EFS ID:	28621844
Application Number:	15061021
International Application Number:	
Confirmation Number:	4734
Title of Invention:	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
First Named Inventor/Applicant Name:	Matthew A. Markley
Customer Number:	98841
Filer:	Paul J. Otterstedt
Filer Authorized By:	
Attorney Docket Number:	YOR920120450US2
Receipt Date:	14-MAR-2017
Filing Date:	04-MAR-2016
Time Stamp:	11:59:01
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	\$ 180
RAM confirmation Number	031417INTEFSW00010393500510
Deposit Account	500510
Authorized User	A. Kammer

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.16 (National application filing, search, and examination fees)

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	YOR920120450US2_IDS_03-14-17.pdf	613545	no	5
			926b1b4f9f022e7cfa7f04a09d11891f2d916684		
Warnings:					
Information:					
2	Foreign Reference	YOR920120450US2_PA1_03-14-17.pdf	2065724	no	35
			ec652d46847f3c632d3fad7c1f5dbc68abe980d3		
Warnings:					
Information:					
3	Non Patent Literature	YOR920120450US2_NPL1_03-14-17.pdf	1323830	no	14
			833d09ba255ddc4daabd7f8eea221272481a07ce		
Warnings:					
Information:					
4	Non Patent Literature	YOR920120450US2_NPL2_03-14-17.pdf	1111087	no	12
			d018bed76c4436134b7b99764a03aec7a1b9eba0		
Warnings:					
Information:					
5	Non Patent Literature	YOR920120450US2_NPL3_03-14-17.pdf	1272247	no	11
			bf596e1b3538952fb102e05d00c29aabfffe8721		
Warnings:					
Information:					
6	Non Patent Literature	YOR920120450US2_NPL4_03-14-17.pdf	789213	no	8
			54199cfbed20b4e62a1f73d16d4d39201e0ce1f8		
Warnings:					
Information:					

7	Non Patent Literature	YOR920120450US2_NPL5_03-1 4-17.pdf	367261 ce7901256db78022200d91749e22828870 8ed901	no	8
---	-----------------------	---------------------------------------	------------------------------------------------------------	----	---

Warnings:**Information:**

8	Non Patent Literature	YOR920120450US2_NPL6_03-1 4-17.pdf	27017 98049a702cba58b2a491e2fc89d62d25634 94902	no	1
---	-----------------------	---------------------------------------	-----------------------------------------------------------	----	---

Warnings:**Information:**

9	Fee Worksheet (SB06)	fee-info.pdf	30604 1c50b254ac95809fd9428129fdb70db5cda 04a09	no	2
---	----------------------	--------------	-----------------------------------------------------------	----	---

Warnings:**Information:**

Total Files Size (in bytes):			7600528
-------------------------------------	--	--	---------

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Doc Code: DIST.E.FILE

6455

Document Description: Electronic Terminal Disclaimer - Filed

U.S. Patent and Trademark Office
Department of Commerce

Electronic Petition Request	TERMINAL DISCLAIMER TO OBIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT
Application Number	15061021
Filing Date	04-Mar-2016
First Named Inventor	Matthew Markley
Attorney Docket Number	YOR920120450US2
Title of Invention	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
<input checked="" type="checkbox"/> Filing of terminal disclaimer does not obviate requirement for response under 37 CFR 1.111 to outstanding Office Action	
<input checked="" type="checkbox"/> This electronic Terminal Disclaimer is not being used for a Joint Research Agreement.	
Owner	Percent Interest
INTERNATIONAL BUSINESS MACHINES CORPORATION	100%
<p>The owner(s) with percent interest listed above in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of prior patent number(s)</p> <p>9282166</p> <p>as the term of said prior patent is presently shortened by any terminal disclaimer. The owner hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the prior patent are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.</p> <p>In making the above disclaimer, the owner does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the prior patent, "as the term of said prior patent is presently shortened by any terminal disclaimer," in the event that said prior patent later:</p> <ul style="list-style-type: none">- expires for failure to pay a maintenance fee;- is held unenforceable;- is found invalid by a court of competent jurisdiction;- is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;- has all claims canceled by a reexamination certificate;- is reissued; or- is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer. <p><input checked="" type="radio"/> Terminal disclaimer fee under 37 CFR 1.20(d) is included with Electronic Terminal Disclaimer request.</p>	

☐ I certify, in accordance with 37 CFR 1.4(d)(4), that the terminal disclaimer fee under 37 CFR 1.20(d) required for this terminal disclaimer has already been paid in the above-identified application.

Applicant claims the following fee status:

- ☐ Small Entity
- ☐ Micro Entity
- ☒ Regular Undiscounted

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

THIS PORTION MUST BE COMPLETED BY THE SIGNATORY OR SIGNATORIES

I certify, in accordance with 37 CFR 1.4(d)(4) that I am:

- ☒ An attorney or agent registered to practice before the Patent and Trademark Office who is of record in this application
- Registration Number 48909
- ☐ A sole inventor
- ☐ A joint inventor; I certify that I am authorized to sign this submission on behalf of all of the inventors as evidenced by the power of attorney in the application
- ☐ A joint inventor; all of whom are signing this request

Signature	/Nathaniel Wallace/
Name	Nathaniel Wallace

*Statement under 37 CFR 3.73(b) is required if terminal disclaimer is signed by the assignee (owner).
Form PTO/SB/96 may be used for making this certification. See MPEP § 324.

Electronic Patent Application Fee Transmittal				
Application Number:		15061021		
Filing Date:		04-Mar-2016		
Title of Invention:		MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION		
First Named Inventor/Applicant Name:		Matthew A. Markley		
Filer:		Nathaniel Thomas Wallace		
Attorney Docket Number:		YOR920120450US2		
Filed as Large Entity				
Filing Fees for Utility under 35 USC 111(a)				
Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Basic Filing:				
STATUTORY OR TERMINAL DISCLAIMER	1814	1	160	160
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Extension-of-Time:				
Miscellaneous:				
Total in USD (\$)				160

Doc Code: DISQ.E.FILE

Document Description: Electronic Terminal Disclaimer – Approved

Application No.: 15061021

Filing Date: 04-Mar-2016

Applicant/Patent under Reexamination: Markley

Electronic Terminal Disclaimer filed on March 14, 2017

☒ APPROVED

This patent is subject to a terminal disclaimer

☐ DISAPPROVED

Approved/Disapproved by: Electronic Terminal Disclaimer automatically approved by EFS-Web

U.S. Patent and Trademark Office

Electronic Acknowledgement Receipt

EFS ID:	28624811
Application Number:	15061021
International Application Number:	
Confirmation Number:	4734
Title of Invention:	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
First Named Inventor/Applicant Name:	Matthew A. Markley
Customer Number:	98841
Filer:	Nathaniel Thomas Wallace
Filer Authorized By:	
Attorney Docket Number:	YOR920120450US2
Receipt Date:	14-MAR-2017
Filing Date:	04-MAR-2016
Time Stamp:	14:59:10
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	\$ 160
RAM confirmation Number	031517INTEFSW00001170500510
Deposit Account	500510
Authorized User	A. Kammer

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:
37 CFR 1.17 (Patent application and reexamination processing fees)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Electronic Terminal Disclaimer-Filed	eTerminal-Disclaimer.pdf	33492	no	2
			a443bbb1c9c17370ec0e46aa36a302cec71cf556		

Warnings:

Information:

2	Fee Worksheet (SB06)	fee-info.pdf	30323	no	2
			4c3984e110d374f1bfafaaa322fc6ff190677dc2		

Warnings:

Information:

Total Files Size (in bytes):			63815
-------------------------------------	--	--	-------

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): Matthew A. Markley
Docket No.: YOR920120450US2 (150-126 Con)
Serial No.: 15/061,021
Filing Date: 03/04/2016
Group: 2453
Examiner: NGUYEN, TU T

Title: MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RESPONSE

Examiner:

In response to the outstanding Office Action dated December 16, 2016, please consider the following remarks.

CLAIMS:

Listing of claims.

1. (Original) A non-transitory computer readable medium comprising computer executable instructions which when executed by a computer cause the computer to perform the method of:

discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component, wherein said at least one source infrastructure management component is an instance of an image, and wherein said at least one source infrastructure management component is running in a customer environment;

querying a database to obtain a description of a target cloud infrastructure;

analyzing said at least one source infrastructure management component using said description of said target cloud infrastructure to determine that said at least one source infrastructure management component is appropriate for infrastructure configuration mapping to said target cloud infrastructure; stopping an application executing on said at least one source infrastructure management component determined appropriate for infrastructure configuration mapping; and

capturing said at least one source infrastructure management component determined appropriate for infrastructure configuration mapping for migration to said target cloud infrastructure.

2. (Original) The non-transitory computer readable medium of Claim 1, wherein, in said discovering step of said method, said at least one source infrastructure management component comprises at least one of:

at least one source infrastructure management client;

at least one source infrastructure management server;

at least one source infrastructure management configuration; and

at least one source infrastructure management log.

3. (Original) The non-transitory computer readable medium of Claim 2, wherein said computer executable instructions further cause the computer to perform the additional method step of deriving a description of non-functional requirements of said source management infrastructure.

4. (Original) The non-transitory computer readable medium of Claim 3, wherein:

in said discovering step of said method, said at least one source infrastructure management component comprises at least:

said at least one source infrastructure management configuration; and

said at least one source infrastructure management log;

in said querying step of said method, said description of said target cloud infrastructure comprises at least one of:

cloud infrastructure software standards;

cloud infrastructure software configurations; and

a description of non-functional requirements of a target management infrastructure of said target cloud infrastructure;

said analyzing step of said method comprises:

matching said cloud infrastructure software standards with at least a corresponding one of:

said at least one source infrastructure management client; and

said at least one source infrastructure management server;

mapping said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure; and

mapping said cloud infrastructure software configurations with said at least one source infrastructure management configuration and said at least one source infrastructure management log.

5. (Original) The non-transitory computer readable medium of Claim 4, wherein said mapping of said cloud infrastructure software configurations with said at least one source infrastructure management configuration and said at least one source infrastructure management log is at least partially based on:

said matching of said cloud infrastructure software standards with said at least one source infrastructure management client and said at least one source infrastructure management server; and

said mapping of said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure.

6. (Original) The non-transitory computer readable medium of Claim 5, wherein:

in said discovering step of said method, said at least one source infrastructure management component further comprises at least one source infrastructure management process;

in said querying step of said method, said description of said target cloud infrastructure further comprises at least one target infrastructure management process; and

said analyzing step of said method further comprises mapping said at least one source infrastructure management process with said at least one target infrastructure management process.

7. (Original) The non-transitory computer readable medium of Claim 6, wherein said mapping of said at least one source infrastructure management process with said at least one target infrastructure management process is at least partially based on said mapping of said description of said non-functional requirements of said source management infrastructure with said description of said non-functional requirements of said target management infrastructure.

8. (Original) The non-transitory computer readable medium of Claim 2, wherein:

in said querying step of said method, said description of said target cloud infrastructure comprises at least one of:

cloud infrastructure software standards; and

cloud infrastructure software configurations;

said analyzing step of said method comprises:

matching said cloud infrastructure software standards with at least a corresponding one of:

said at least one source infrastructure management client; and

said at least one source infrastructure management server; and

mapping said cloud infrastructure software configurations with said at least one source infrastructure management configuration.

9. (Original) The non-transitory computer readable medium of Claim 8, wherein said mapping of said cloud infrastructure software configurations with said at least one source infrastructure management configuration is at least partially based on said matching of said cloud infrastructure software standards with said at least one source infrastructure management client and said at least one source infrastructure management server.

10. (Original) The non-transitory computer readable medium of Claim 9, wherein:

in said discovering step of said method, said at least one source infrastructure management component further comprises at least one source infrastructure management process;

in said querying step of said method, said description of said target cloud infrastructure further comprises at least one target infrastructure management process; and

said analyzing step of said method further comprises mapping said at least one source infrastructure management process with said at least one target infrastructure management process.

11. (Original) The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method further comprises determining that at least one conflict exists with said at least one mandatory target infrastructure management component, wherein said at least one conflict comprises said at least one source infrastructure management component managing at least one object that said at least one mandatory target infrastructure management component will manage in said target cloud infrastructure.

12. (Original) The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method further comprises determining that at least one conflict exists with said at least one mandatory target infrastructure management component, wherein said at least one conflict comprises said at least one source infrastructure management component using at least one resource that said at least one mandatory target infrastructure management component will use in said target cloud infrastructure.

13. (Original) The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method further comprises determining that at least one conflict exists with said at least one mandatory target infrastructure management component, wherein said at least one conflict comprises at least one of:

- current lack of a mandatory target infrastructure component on a client,
- existence of a mandatory target infrastructure component in a different version, and
- existence of a mandatory target infrastructure component with different configurations.

14. (Original) The non-transitory computer readable medium of Claim 1, wherein said analyzing step of said method comprises at least one of:

- recommending un-installation of a source infrastructure component;

recommending installation of a target infrastructure component;
recommending modifying configuration of an infrastructure component; and
recommending excluding servers with a given source infrastructure component from migration.

15. (Original) The non-transitory computer readable medium of Claim 1, wherein:

said discovering step of said method comprises discovering, in said source computing system, a plurality of source infrastructure management components; and

said analyzing step of said method comprises:

initially aggregating all of said plurality of source infrastructure management components;

developing generic recommendations for said aggregated plurality of source infrastructure management components;

applying said recommendations on a per-instance basis; and

refining at least some of said applied recommendations in cases where said generic recommendations are inconclusive.

16. (Original) The non-transitory computer readable medium of Claim 1, wherein said discovering step of said method is carried out on at least one physical machine in said source computing system, said at least one physical machine comprising a client of said at least one source infrastructure management component.

17. (Original) The non-transitory computer readable medium of Claim 1, wherein said discovering step of said method is carried out on at least one physical machine in said source computing system, said at least one physical machine comprising a server on which said at least one source infrastructure management component runs.

18. (Original) The non-transitory computer readable medium of Claim 1, wherein said computer executable instructions comprise distinct software modules, each of the distinct software modules being embodied on said non-transitory computer readable medium, and wherein the distinct software modules comprise a discovery tool module, a description module, and an infrastructure comparison engine module;

wherein:

said discovering is carried out by said discovery tool module executing on said computer;
said querying is carried out by said description module executing on said computer; and
said analyzing is carried out by said infrastructure comparison engine module executing on said computer.

19. (Original) A non-transitory computer readable medium comprising computer executable instructions which when executed by a computer cause the computer to perform the method of:

discovering, in a source computing system having a source management infrastructure, at least one source infrastructure management component, wherein said at least one source infrastructure management component is an instance of an image, and wherein said at least one source infrastructure management component is running in a customer environment;

querying a database to obtain a description of a target cloud infrastructure;

analyzing said at least one source infrastructure management component using said description of said target cloud infrastructure to determine that said at least one source infrastructure management component is appropriate for infrastructure configuration mapping to said target cloud infrastructure, wherein:

in said discovering step, said at least one source infrastructure management component comprises at least one of:

at least one source infrastructure management client;

at least one source infrastructure management server;

at least one source infrastructure management configuration; and

at least one source infrastructure management log;

in said querying step, said description of said target cloud infrastructure comprises at least one of:

cloud infrastructure software standards; and

cloud infrastructure software configurations;

said analyzing comprises:

matching said cloud infrastructure software standards with at least a corresponding one

of:

said at least one source infrastructure management client; and

said at least one source infrastructure management server; and

mapping said cloud infrastructure software configurations with said at least one source infrastructure management configuration determined appropriate for infrastructure configuration mapping.

REMARKS

The present application was filed on March 4, 2016, with Claims 1-19. Claims 1-19 are presently pending in the above-identified patent application.

In the Office Action the Examiner:

1) rejected Claim 1-19 on the ground of nonstatutory double patenting as being unpatentable over Claims 1-19 of U.S. Patent No. 9,282,166; and

2) indicated that Claims 1-19 would be allowable upon filing of a Terminal Disclaimer to overcome the Double Patenting rejection.

In response to the Office Action, Applicants have filed an Electronic Terminal Disclaimer dated March 14, 2017, together with the terminal disclaimer fee under 37 CFR 1.20(d), where the Electronic Terminal Disclaimer references U.S. Patent No. 9,282,166. Applicants respectfully request withdrawal of the rejection in view of the Terminal Disclaimer.

Conclusion

All of the pending claims, i.e., Claims 1-19, are in condition for allowance and such favorable action is earnestly solicited.

If any outstanding issues remain, or if the Examiner has any further suggestions for expediting allowance of this application, the Examiner is invited to contact the undersigned at the telephone number indicated below.

The Examiner's attention to this matter is appreciated.

Respectfully submitted,

Dated: March 14, 2017

By: /Nathaniel T. Wallace/
Nathaniel T. Wallace
Reg. No. 48,909
Paul J. Otterstedt
Reg. No. 37,411
Attorneys for Applicant

Otterstedt, Ellenbogen & Kammer, LLP
P.O. Box 381
Cos Cob, CT 06807-0381
(203) 433-3333

Electronic Acknowledgement Receipt

EFS ID:	28624869
Application Number:	15061021
International Application Number:	
Confirmation Number:	4734
Title of Invention:	MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION
First Named Inventor/Applicant Name:	Matthew A. Markley
Customer Number:	98841
Filer:	Nathaniel Thomas Wallace
Filer Authorized By:	
Attorney Docket Number:	YOR920120450US2
Receipt Date:	14-MAR-2017
Filing Date:	04-MAR-2016
Time Stamp:	15:01:12
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		YOR920120450US2_Respons oOfficeAction1_03-14-17.pdf	96536	yes	9
			8b9e111c2529e11ad5d4aed514dfbf8e474 cc3ab		

Multipart Description/PDF files in .zip description			
Document Description		Start	End
Amendment/Req. Reconsideration-After Non-Final Reject		1	1
Claims		2	8
Amendment/Req. Reconsideration-After Non-Final Reject		9	9

Warnings:**Information:**

Total Files Size (in bytes):	96536
-------------------------------------	-------

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875					Application or Docket Number 15/061,021		Filing Date 03/04/2016		<input type="checkbox"/> To be Mailed		
ENTITY: <input checked="" type="checkbox"/> LARGE <input type="checkbox"/> SMALL <input type="checkbox"/> MICRO											
APPLICATION AS FILED – PART I											
(Column 1)			(Column 2)								
FOR		NUMBER FILED		NUMBER EXTRA		RATE (\$)		FEE (\$)			
<input type="checkbox"/> BASIC FEE (37 CFR 1.16(a), (b), or (c))		N/A		N/A		N/A					
<input type="checkbox"/> SEARCH FEE (37 CFR 1.16(k), (i), or (m))		N/A		N/A		N/A					
<input type="checkbox"/> EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))		N/A		N/A		N/A					
TOTAL CLAIMS (37 CFR 1.16(i))		minus 20 =		*		X \$ =					
INDEPENDENT CLAIMS (37 CFR 1.16(h))		minus 3 =		*		X \$ =					
<input type="checkbox"/> APPLICATION SIZE FEE (37 CFR 1.16(s))		If the specification and drawings exceed 100 sheets of paper, the application size fee due is \$310 (\$155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).									
<input type="checkbox"/> MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))											
* If the difference in column 1 is less than zero, enter "0" in column 2.						TOTAL					
APPLICATION AS AMENDED – PART II											
(Column 1)			(Column 2)			(Column 3)					
AMENDMENT	03/14/2017		CLAIMS REMAINING AFTER AMENDMENT			HIGHEST NUMBER PREVIOUSLY PAID FOR		PRESENT EXTRA			
	Total (37 CFR 1.16(i))		* 19		Minus	** 20		= 0			
	Independent (37 CFR 1.16(h))		* 2		Minus	*** 3		= 0			
	<input type="checkbox"/> Application Size Fee (37 CFR 1.16(s))										
	<input type="checkbox"/> FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))										
						TOTAL ADD'L FEE		0			
(Column 1)			(Column 2)			(Column 3)					
AMENDMENT			CLAIMS REMAINING AFTER AMENDMENT			HIGHEST NUMBER PREVIOUSLY PAID FOR		PRESENT EXTRA			
	Total (37 CFR 1.16(i))		*		Minus	**		=			
	Independent (37 CFR 1.16(h))		*		Minus	***		=			
	<input type="checkbox"/> Application Size Fee (37 CFR 1.16(s))										
	<input type="checkbox"/> FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))										
						TOTAL ADD'L FEE					
<p>* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.</p> <p>** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".</p> <p>*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".</p> <p>The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.</p>											

LIE
MARQUITA JONES

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

98841 7590 03/24/2017
 Otterstedt, Ellenbogen & Kammer, LLP
 P.O. Box 381
 Cos Cob, CT 06807-0381

EXAMINER

NGUYEN, TU T

ART UNIT

PAPER NUMBER

2453

DATE MAILED: 03/24/2017

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

15/061,021

03/04/2016

Matthew A. Markley

YOR920120450US2

4734

TITLE OF INVENTION: MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	UNDISCOUNTED	\$960	\$0	\$0	\$960	06/26/2017

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PART B - FEE(S) TRANSMITTAL

6475

Complete and send this form, together with applicable fee(s), to: **Mail****Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (571)-273-2885**

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

98841 7590 03/24/2017
Otterstedt, Ellenbogen & Kammer, LLP
P.O. Box 381
Cos Cob, CT 06807-0381

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/061,021	03/04/2016	Matthew A. Markley	YOR920120450US2	4734

TITLE OF INVENTION: MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	UNDISCOUNTED	\$960	\$0	\$0	\$960	06/26/2017

EXAMINER	ART UNIT	CLASS-SUBCLASS
NGUYEN, TU T	2453	709-203000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

- ☐ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
- ☐ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list

- (1) The names of up to 3 registered patent attorneys or agents OR, alternatively, 1 _____
- (2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed. 2 _____
- 3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE

(B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent): ☐ Individual ☐ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:

- ☐ Issue Fee
- ☐ Publication Fee (No small entity discount permitted)
- ☐ Advance Order - # of Copies _____

4b. Payment of Fee(s): (Please first reapply any previously paid issue fee shown above)

- ☐ A check is enclosed.
- ☐ Payment by credit card. Form PTO-2038 is attached.
- ☐ The director is hereby authorized to charge the required fee(s), any deficiency, or credits any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).

5. Change in Entity Status (from status indicated above)

- ☐ Applicant certifying micro entity status. See 37 CFR 1.29
- ☐ Applicant asserting small entity status. See 37 CFR 1.27
- ☐ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.

NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.

NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____

Date _____

Typed or printed name _____

Registration No. _____



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/061,021	03/04/2016	Matthew A. Markley	YOR920120450US2	4734

98841	7590	03/24/2017
-------	------	------------

Otterstedt, Ellenbogen & Kammer, LLP
 P.O. Box 381
 Cos Cob, CT 06807-0381

EXAMINER
NGUYEN, TU T

ART UNIT	PAPER NUMBER
2453	

DATE MAILED: 03/24/2017

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
 (Applications filed on or after May 29, 2000)

The Office has discontinued providing a Patent Term Adjustment (PTA) calculation with the Notice of Allowance.

Section 1(h)(2) of the AIA Technical Corrections Act amended 35 U.S.C. 154(b)(3)(B)(i) to eliminate the requirement that the Office provide a patent term adjustment determination with the notice of allowance. See Revisions to Patent Term Adjustment, 78 Fed. Reg. 19416, 19417 (Apr. 1, 2013). Therefore, the Office is no longer providing an initial patent term adjustment determination with the notice of allowance. The Office will continue to provide a patent term adjustment determination with the Issue Notification Letter that is mailed to applicant approximately three weeks prior to the issue date of the patent, and will include the patent term adjustment on the patent. Any request for reconsideration of the patent term adjustment determination (or reinstatement of patent term adjustment) should follow the process outlined in 37 CFR 1.705.

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 15/061,021	Applicant(s) MARKLEY ET AL.	
	Examiner Tu Nguyen	Art Unit 2453	AIA (First Inventor to File) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to 03/14/2017.
☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on ____.
2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on ____; the restriction requirement and election have been incorporated into this action.
3. ☒ The allowed claim(s) is/are 1-19. As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.
4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a) ☐ All b) ☐ Some *c) ☐ None of the:

1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. ____.
3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

* Certified copies not received: ____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date ____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

<ol style="list-style-type: none"> 1. <input type="checkbox"/> Notice of References Cited (PTO-892) 2. <input checked="" type="checkbox"/> Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date ____ 3. <input type="checkbox"/> Examiner's Comment Regarding Requirement for Deposit of Biological Material 4. <input type="checkbox"/> Interview Summary (PTO-413), Paper No./Mail Date ____. 	<ol style="list-style-type: none"> 5. <input type="checkbox"/> Examiner's Amendment/Comment 6. <input checked="" type="checkbox"/> Examiner's Statement of Reasons for Allowance 7. <input type="checkbox"/> Other ____.
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

/Tu Nguyen/ Primary Examiner, Art Unit 2453	
------------------------------------------------	--

Application/Control Number: 15/061,021
Art Unit: 2453

Page 2

The present application is being examined under the pre-AIA first to invent provisions.

REASONS FOR ALLOWANCE

The following is an examiner's statement of reasons for allowance:

Claims 1-19 are allowed in light of the amendments and the Applicant Arguments/Remarks made in the amendment filed on 03/14/2017.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Tu Nguyen whose telephone number is (571)272-2424. The examiner can normally be reached on M-W,F 7:30-5:30.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at <http://www.uspto.gov/interviewpractice>.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kamal Divecha can be reached on (571) 272-5863. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Application/Control Number: 15/061,021
Art Unit: 2453

Page 3

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Tu Nguyen/
Primary Examiner, Art Unit 2453
03/20/2017

Doc code: IDS

#: 6481

PTO/SB/08a (01-10)

Doc description: Information Disclosure Statement (IDS) Filed

Approved for use through 07/31/2012. OMB 0651-0031

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	15061021
	Filing Date	2016-03-04
	First Named Inventor	Matthew A. Markley
	Art Unit	2453
	Examiner Name	NGUYEN, TU T
	Attorney Docket Number	YOR920120450US2

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8566824		2013-10-22	Traut; Eric P.	
	2	8689054		2014-04-01	van der Goot; Roel	
	3	8726083		2014-05-13	van der Goot; Roel	
If you wish to add additional U.S. Patent citation information please click the Add button.						Add
U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20140089658		2014-03-27	Raghuram; Yeluri ; et al.	
	2	20110295999		2011-12-01	Ferris; James Michael ; et al.	
	3	20130275596		2013-10-17	Subramaniam; Karthikeyan	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number # 6482		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

4	20140059539	2014-02-27	Simonsen; Harold C. ; et al.
5	20110071983	2011-03-24	Murase; Atsushi
6	20130262390	2013-10-03	Kumarasamy; Paramasivam ; et al.
7	20120084775	2012-04-05	Lotlikar; Mahesh ; et al.
8	20100106885	2010-04-29	Gao; Zhi Guo ; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	CN102387197	CN		2012-03-21	Lotlikar; Mahesh ; et al.		

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Clark et al., "Live Migration of Virtual Machines," 2nd Symposium on Networked Systems Design & Implementation, NSDI'05, 2005, pp. 273-286.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number # 6483		15061021
Filing Date		2016-03-04
First Named Inventor	Matthew A. Markley	
Art Unit	2453	
Examiner Name	NGUYEN, TU T	
Attorney Docket Number	YOR920120450US2	

2	Nicolae, Bogdan and Cappello, Franck, "BlobCR: Efficient Checkpoint-Restart for HPC Applications on IaaS Clouds using Virtual Disk Snapshots," Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC'11, 12-18 November 2011, p. 1-12.
3	Park et al., "Fast and Space-Efficient Virtual Machine Checkpointing," Proceedings of the 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE'11, 9-11 March 2011, pp. 75-86.
4	Zhao, Ming and Figueiredo, Renato J., "Experimental Study of Virtual Machine Migration in Support of Reservation of Cluster Resources," Proceedings of the 2nd International Workshop on Virtualization Technology in Distributed Computing, VTDC'07, 12 November 2007, p 1-8.
5	state intellectual property office of the people's republic of China, First Office Action, Chinese Application 201310626100.0, June 28, 2016, includes machine translation, pages 1-8
6	state intellectual property office of the people's republic of China, First Search, Chinese Application 201310626100.0, June 17 2016, page 1

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/TU T NGUYEN/	Date Considered	03/20/2017
--------------------	---------------	-----------------	------------

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number # 6484		15061021	
Filing Date		2016-03-04	
First Named Inventor	Matthew A. Markley		
Art Unit	2453		
Examiner Name	NGUYEN, TU T		
Attorney Docket Number	YOR920120450US2		

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

- ☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.
- ☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Paul J.Otterstedt/	Date (YYYY-MM-DD)	2017-03-14
Name/Print	Paul J.Otterstedt	Registration Number	37411

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.




UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

BIB DATA SHEET

CONFIRMATION NO. 4734

SERIAL NUMBER	FILING or 371(c) DATE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.	
15/061,021	03/04/2016	709	2453	YOR920120450US2	
RULE					
APPLICANTS International Business Machines Corporation, Armonk, NY; INVENTORS Matthew A. Markley, Hays, KS; Amitkumar M. Paradkar, Mohegan Lake, NY; Venkata Vinay Parisa, Charlotte, NC; Birgit M. Pfitzmann, Samstagern, SWITZERLAND; ** CONTINUING DATA ***** This application is a CON of 13/689,500 11/29/2012 PAT 9282166 ** FOREIGN APPLICATIONS ***** ** IF REQUIRED, FOREIGN FILING LICENSE GRANTED ** 03/18/2016					
Foreign Priority claimed <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No 35 USC 119(a-d) conditions met <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No Verified and Acknowledged <u>/TU T NGUYEN/</u> Examiner's Signature	<input type="checkbox"/> Met after Allowance Initials	STATE OR COUNTRY KS	SHEETS DRAWINGS 47	TOTAL CLAIMS 19	INDEPENDENT CLAIMS 2
ADDRESS Otterstedt, Ellenbogen & Kammer, LLP P.O. Box 381 Cos Cob, CT 06807-0381 UNITED STATES					
TITLE MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION					
FILING FEE RECEIVED 2000	FEES: Authority has been given in Paper No. _____ to charge/credit DEPOSIT ACCOUNT No. _____ for following:		<input type="checkbox"/> All Fees <input type="checkbox"/> 1.16 Fees (Filing) <input type="checkbox"/> 1.17 Fees (Processing Ext. of time) <input type="checkbox"/> 1.18 Fees (Issue) <input type="checkbox"/> Other _____ <input type="checkbox"/> Credit		

Search Notes 	Application/Control No. 15061021	Applicant(s)/Patent Under Reexamination MARKLEY ET AL.
	Examiner TU NGUYEN	Art Unit 2453

CPC- SEARCHED		
Symbol	Date	Examiner
G06F 15/16, 15/173, 15/177, 9/455	03/20/2017	TN
H04L 29/08, 67/42, 29/06	03/20/2017	TN


CPC COMBINATION SETS - SEARCHED		
Symbol	Date	Examiner

US CLASSIFICATION SEARCHED			
Class	Subclass	Date	Examiner
709	203	03/20/2017	TN

SEARCH NOTES		
Search Notes	Date	Examiner
East	12/11/2016	TN
709/203, 218-226	12/11/2016	TN
G06F 15/16, 15/173, 15/177, 9/455	12/11/2016	TN
H04L 29/08, 67/42, 29/06	12/11/2016	TN
Search updated	03/20/2017	TN
Inventor search	03/20/2017	TN

INTERFERENCE SEARCH			
US Class/ CPC Symbol	US Subclass / CPC Group	Date	Examiner
	See search history (L47)	03/20/2017	TN


--	--

Issue Classification 	Application/Control No. 15061021	Applicant(s)/Patent Under Reexamination MARKLEY ET AL.	
	Examiner TU NGUYEN	Art Unit 2453	

CPC						
Symbol					Type	Version
H04L		41	/	04	F	2013-01-01
H04L		67	/	34	I	2013-01-01
H04L		67	/	1002	I	2013-01-01
H04W		4	/	003	I	2013-01-01
G06F		17	/	30914	I	2013-01-01
G06F		17	/	30979	I	2013-01-01
H04L		67	/	10	I	2013-01-01
H04L		41	/	12	I	2013-01-01
			/			
			/			
			/			
			/			
			/			
			/			
			/			


CPC Combination Sets				
Symbol	Type	Set	Ranking	Version

NONE		Total Claims Allowed:	
		19	
(Assistant Examiner)	(Date)	O.G. Print Claim(s)	O.G. Print Figure
/TU NGUYEN/ Primary Examiner.Art Unit 2453	03/20/2017	1	5
(Primary Examiner)	(Date)		

Issue Classification 	Application/Control No. 15061021	Applicant(s)/Patent Under Reexamination MARKLEY ET AL.
	Examiner TU NGUYEN	Art Unit 2453

[illegible]

NONE		Total Claims Allowed:	
(Assistant Examiner)	(Date)	19	
/TU NGUYEN/ Primary Examiner.Art Unit 2453	03/20/2017	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	5

Issue Classification 	Application/Control No. 15061021	Applicant(s)/Patent Under Reexamination MARKLEY ET AL.
	Examiner TU NGUYEN	Art Unit 2453

<input checked="" type="checkbox"/> Claims renumbered in the same order as presented by applicant																<input type="checkbox"/> CPA		<input checked="" type="checkbox"/> T.D.		<input type="checkbox"/> R.1.47	
Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original						
	1		17																		
	2		18																		
	3		19																		
	4																				
	5																				
	6																				
	7																				
	8																				
	9																				
	10																				
	11																				
	12																				
	13																				
	14																				
	15																				
	16																				

NONE		Total Claims Allowed:	
(Assistant Examiner)		19	
(Date)			
/TU NGUYEN/ Primary Examiner.Art Unit 2453		03/20/2017	
(Primary Examiner)		(Date)	
		O.G. Print Claim(s)	O.G. Print Figure
		1	5

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	48	("20060218544" "20100287280" "20100306772" "20110060832" "20110202714" "20120072903" "20120304179" "20130091285" "20140146055" "20140149492" "20140149494" "20140149591" "7844853" "7191441" "7356679" "20120131567" "20110320882" "20110107331" "20090307166" "20090063751" "20100077160" "20120117558" "20140129521" "20090288084" "20130036091" "20070033356" "20140115285" "20140149695" "20120266169" "20130167145" "20100049929" "20120209812" "20100191930" "20140149983" "20150331708" "9282166" "9092837" "9292330" "20050204040" "20120203823" "20120311154" "20120137003" "20080189697" "20100235831" "20110106768" "20090249284" "20160110215" "20160191332").PN.	US-PGPUB; USPAT	OR	OFF	2017/03/20 09:42
L2	20	((("7844853") or ("7191441") or ("7356679") or ("20120131567") or ("20110320882") or ("20110107331") or ("20090307166") or ("20090063751") or ("20100077160") or ("20120117558") or ("20140129521") or ("20090288084") or ("20130036091") or ("20070033356") or ("20140115285") or ("20140149695") or ("20120266169") or ("20130167145") or ("20100049929") or ("20120209812"))).PN.	US-PGPUB; USPAT	OR	OFF	2017/03/20 09:42
L3	15033	cloud with ((management or manager) near3 (component or module or system or apparatus))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L4	305	target with cloud with infra\$1structure	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L5	46	L3 and L4 and @ad<="20122911"	US-PGPUB; USPAT;	ADJ	ON	2017/03/20 09:42

6492

			USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			
L6	159318	(map or mapping) near7 \$5function\$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L7	6250266	configuration	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L8	319	L3 and L6 and L7 and @ad<="20122911"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L9	25803	cloud with infra\$1structure	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L10	242	L8 and L9	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L11	3419209	target	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L12	121	L10 and L11	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L13	114479	query\$5 near3 database	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42

L14	24	L12 and L13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L15	3	(L1 or L2) and L4	USPAT	OR	OFF	2017/03/20 09:42
L16	4588	cloud with migrat\$5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L17	183	L16 and L6 and L9 and @ad<="20121129"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L18	13	L17 and L13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L19	533	((program or application) near10 (run or running or execut\$5)) with ((map or mapping) near7 configuration)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L20	1	L17 and L19	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:42
L22	11	("8566824" "8689054" "8726083" "20140089658" "20110295999" "20130275596" "20140059539" "20110071983" "20130262390" "20120084775" "20100106885").PN.	US-PGPUB; USPAT	OR	OFF	2017/03/20 09:42
L25	47	17 and 3	US-PGPUB; USPAT	OR	OFF	2017/03/20 09:57
L26	15033	cloud with ((management or manager) near3 (component or module or system or apparatus))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L27	159318	(map or mapping) near7 \$5function\$4	US-PGPUB; USPAT;	ADJ	ON	2017/03/20 09:58

6494

			USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			
L28	6250266	configuration	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L29	319	L26 and L27 and L28 and @ad<="20122911"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L30	25803	cloud with infra\$1structure	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L31	242	L29 and L30	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L32	3419209	target	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L33	121	L31 and L32	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L34	136078	(709/203,218-226).CCLS.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2017/03/20 09:58
L35	19568	G06F15/16.cpc. or G06F15/173.cpc. or G06F15/177.cpc. or G06F9/455.cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58

L36	105465	H04L29/08.cpc. or H04L67/42.cpc. or H04L29/06.cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L37	244867	L34 or L35 or L36	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L38	41	L37 and L33	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 09:58
L40	18	((Matthew) near2 (Markley)).INV.	US-PGPUB; USPAT; USOCR	ADJ	ON	2017/03/20 10:04
L41	34	((Amitkumar) near2 (Paradkar)).INV.	US-PGPUB; USPAT; USOCR	ADJ	ON	2017/03/20 10:04
L42	8	((Venkata) near2 (Parisa)).INV.	US-PGPUB; USPAT; USOCR	ADJ	ON	2017/03/20 10:04
L43	108	((Birgit) near2 (Pitzmann)).INV.	US-PGPUB; USPAT; USOCR	ADJ	ON	2017/03/20 10:04
L44	130	L40 or L41 or L42 or L43	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 10:04
L45	14	L44 and L26	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 10:04
L47	10	((cloud with ((management or manager) near3 (component or module or system or apparatus))) and ((map or mapping) near7 \$5function\$4)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2017/03/20 10:05

3/ 20/ 2017 10:08:27 AM

C:\Users\nguyen24\Documents\EAST\Workspaces\15061021.wsp

PTO/AIA/122 (08-12)

Approved for use through 01/31/2018. OMB 0651-0035

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**CHANGE OF
CORRESPONDENCE ADDRESS**
*Application*Address to:
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Application Number	15/061,021
Filing Date	03-04-2016
First Named Inventor	Matthew A. Markley
Art Unit	2453
Examiner Name	NGUYEN, TU T
Attorney Docket Number	YOR920120450US2

Please change the Correspondence Address for the above-identified patent application to:

☒ The address associated with
Customer Number:

00877

OR

☐ Firm or
Individual Name

Address

City

State

Zip

Country

Telephone

Email

This form cannot be used to change the data associated with a Customer Number. To change the data associated with an existing Customer Number use "Request for Customer Number Data Change" (PTO/SB/124).

I am the:

☐

Applicant

☒

Attorney or agent of record. Registration Number 37,411

☐

Registered practitioner named in the application transmittal papers who acts in a representative capacity under 37 CFR 1.34. See 37 CFR 1.33(a)(1). Registration Number

Signature

Typed or Printed
Name

PAUL J. OTTERSTEDT

Date 16 JUNE 2017

Telephone 203-433-3157

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications.
Submit multiple forms if more than one signature is required, see below.☒

*Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: **Mail** Mail Stop ISSUE FEE
 Commissioner for Patents
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
or Fax (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

98841 7590 03/24/2017
 Otterstedt, Ellenbogen & Kammer, LLP
 P.O. Box 381
 Cos Cob, CT 06807-0381



Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/061,021	03/04/2016	Matthew A. Markley	YOR920120450US2	4734

TITLE OF INVENTION: MANAGEMENT INFRASTRUCTURE ANALYSIS FOR CLOUD MIGRATION

APPL. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	UNDISCOUNTED	\$960	\$0	\$0	\$960	06/26/2017

EXAMINER	ART UNIT	CLASS-SUBCLASS
NGUYEN, TU T	2453	709-203000

06/23/2017 CCHAU2 00000025 500510 15061021
 01 FC:1501 960.00 DA

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

☒ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

☐ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. Use of a Customer Number is required.

2. For printing on the patent front page, list

(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,

(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

Louis J. Percello

1 Otterstedt, Ellenbogen & Kammer, LLP

2

3

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE

(B) RESIDENCE: (CITY and STATE OR COUNTRY)

INTERNATIONAL BUSINESS MACHINES CORPORATION ARMONK, NEW YORK

Please check the appropriate assignee category or categories (will not be printed on the patent): ☐ Individual ☒ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:

- ☒ Issue Fee
☒ Publication Fee (No small entity discount permitted)
☐ Advance Order - # of Copies _____

4b. Payment of Fee(s): (Please first reapply any previously paid issue fee shown above)

- ☐ A check is enclosed.
☐ Payment by credit card. Form PTO-2038 is attached.
☒ The director is hereby authorized to charge the required fee(s), any deficiency, or credits any overpayment, to Deposit Account Number 50-0510 (enclose an extra copy of this form).

5. Change in Entity Status (from status indicated above)

- ☐ Applicant certifying micro entity status. See 37 CFR 1.29
☐ Applicant asserting small entity status. See 37 CFR 1.27
☐ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.

NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.

NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature

PAUL J. OTTERSTEDT

Date

JUNE 17, 2017

Typed or printed name

Registration No.

37411



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/061,021	08/01/2017	9722858	YOR920120450US2	4734

877 7590 07/12/2017
IBM CORPORATION
IP Law Department
T.J. Watson Research Center
1101 Kitchawan Road, Route 134, P.O. Box 218
Yorktown, NY 10598

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b) (application filed on or after May 29, 2000)

The Patent Term Adjustment is 0 day(s). Any patent to issue from the above-identified application will include an indication of the adjustment on the front page.

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (<http://pair.uspto.gov>).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Application Assistance Unit (AAU) of the Office of Data Management (ODM) at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site <http://pair.uspto.gov> for additional applicants):

Matthew A. Markley, Hays, KS;
International Business Machines Corporation, Armonk, NY;
Amitkumar M. Paradkar, Mohegan Lake, NY;
Venkata Vinay Parisa, Charlotte, NC;
Birgit M. Pfitzmann, Samstagern, SWITZERLAND;

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The USA offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to encourage and facilitate business investment. To learn more about why the USA is the best country in the world to develop technology, manufacture products, and grow your business, visit SelectUSA.gov.